arXiv:1603.03977v2 [cs.LG] 23 Aug 2016

Privacy-Preserving Analysis of Correlated Data

Yizhen Wang
UC San Diego
yiw248@eng.ucsd.edu

Shuang Song
UC San Diego
shs037@eng.ucsd.edu

Kamalika Chaudhuri
UC San Diego
kamalika@eng.ucsd.edu

Abstract

Many modern machine learning applications involve sensitive correlated data, such as private information on users connected together in a social network, and measurements of physical activity of a single user across time. However, the current standard of privacy in machine learning, differential privacy, does not adequately address privacy issues in this kind of data.

This work looks at a recent generalization of differential privacy, called Pufferfish, that can be used to address privacy in correlated data. The main challenge in applying Pufferfish is a lack of suitable mechanisms. In this paper, we provide a general mechanism, called the Wasserstein Mechanism, which applies to any Pufferfish framework. Since the Wasserstein Mechanism may be computationally inefficient, we provide an additional mechanism that applies to some practical cases such as physical activity measurements across time, is computationally efficient, and provides both privacy and utility.

1 Introduction

Machine-learning is increasingly done on personal data, and consequently it is important to design methods that can learn while still preserving privacy of individuals in sensitive datasets. For the past several years, the de facto standard of privacy in machine learning has been differential privacy [7], and there is a growing body of work on differentially private learning algorithms [4, 15, 17, 5, 12, 24, 6, 2, 14]. The typical setting in these works is that each (independently drawn) training example is a single individual's private value, and the goal is to learn a model while adding enough noise to hide the evidence of the participation of a single individual in the dataset.

Machine learning applications also increasingly involve a different setting – correlated data – privacy issues in which have been largely ignored by this literature. Consider, for example, a simple application – measurement of physical activity of a single subject across time. The goal here is to release aggregate statistics on the subject's activities over a long period (say a week) while hiding activity at any specific instant (say, 10:30am on Jan 4). If the measurements are made at small intervals, then they are highly correlated as human activities change slowly over time. While differential privacy does not directly apply to this example, a modified version called *entry-privacy* [12] is applicable. Entry-privacy ensures that the inclusion of a single time-series entry does not affect the probability of any outcome by much and will advocate adding noise with standard deviation ≈ 1 to each bin of the histogram of activities. While this is enough for independent activities, it is insufficient to hide the evidence of an activity at a specific instant as this activity is likely to continue for several instants. A related notion is group privacy [10], which extends the definition to participation of *entire groups* of correlated individuals or entries. In this example, however, all entries are correlated, and hence group privacy will add $\approx \mathcal{O}(T)$ noise to a histogram over T measurements, which will destroy all utility. Thus to address privacy challenges in this kind of data, we need a different notion of privacy.

A generalized version of differential privacy called Pufferfish that addresses some of these issues was proposed by [19]. In Pufferfish, privacy requirements are specified through three components – $\mathcal{S}$, a set of secrets that represents what may need to be hidden, $\mathcal{Q}$, a set of secret pairs that represents

pairs of secrets that need to be indistinguishable to the adversary and Θ, a class of distributions that can plausibly generate the data. Privacy is provided by ensuring that the secret pairs in $\mathcal{Q}$ are indistinguishable when data is generated from any $\theta \in \Theta$. In our physical activity example, $\mathcal{S}$ is the set of activities at each time t, and secret pairs are all pairs of the form (Activity a at time t, Activity b at time t). Assuming that activities transition in a Markovian fashion, Θ is a set of Markov Chains over activities. Pufferfish captures correlation in such applications effectively in two ways – first, unlike differential privacy, it can hide private values against correlation across multiple entries/individuals; second, unlike group privacy, it also allows utility in cases where a large number of individuals or entries are correlated, yet the "average amount" of correlation is low. Because of these properties, we adopt Pufferfish as our privacy definition.

The main challenge in using Pufferfish is a lack of suitable mechanisms. While mechanisms are known for specific frameworks [19, 13, 25], there is no general mechanism, and moreover, there is no existing mechanism that applies to the physical activity measurement example. Our first contribution is to provide a *generic mechanism* that provides privacy in *any Pufferfish framework*. Our mechanism, called the Wasserstein Mechanism, is an analogue of the popular global sensitivity mechanism for differential privacy, and exploits a novel connection between privacy and Wasserstein distance between measures. However, this mechanism, due to its extreme generality, can be computationally inefficient. We next turn our attention to the case when correlation between variables can be described by a Bayesian network, and the goal is to hide the private value of each variable. We provide a second mechanism, called the Markov Quilt Mechanism, that can exploit properties of the Bayesian network to reduce the computational complexity. As a case study, we apply our mechanism to a Markov Chain, which models the physical activity monitoring problem, and we show that it runs in $\mathcal{O}(T^3)$ time for chains of size T. We also show that when Θ consists of rapidly mixing chains, the resulting mechanism adds noise with standard deviation $\approx \mathcal{O}(1)$ to each bin of the activity histogram, and thus it offers both privacy and utility.

Related Work. There is a great deal of work on differentially private mechanisms – see [23, 10]. As we explain earlier, however, differential privacy provides insufficient privacy for the kind of applications we consider. A related framework is coupled-worlds privacy [1]; while this framework can take data distributions into account through a distribution family Θ, it too requires that the participation of a single individual or entry in the data does not make a difference, and thus it provides insufficient privacy for our applications. We remark that while mechanisms for specific coupled-worlds privacy frameworks exist, there is also no *generic* coupled-worlds privacy mechanism.

Our work instead uses Pufferfish, a recent generalization of differential privacy [19]. [19, 13] provide some examples of Pufferfish frameworks with associated privacy mechanisms; however, *they do not provide a generic Pufferfish mechanism, and none of their examples apply to Bayesian networks.* [25] uses a modification of Pufferfish, where instead of a distribution class Θ, they consider a single generating distribution θ. However, *instead of generic mechanisms*, they consider correlations that can be modeled by Gaussian Markov Random Fields, and thus their work also does not apply to our physical activity monitoring example. [16] apply Pufferfish to smart-meter data; instead of directly modeling correlation through Markov models, they add noise to the wavelet coefficients of the time series that correspond to different frequencies.

In concurrent work,[1] [11] provide an alternative algorithm for Pufferfish privacy when the data X can be written as $X = (X_1, \ldots, X_n)$ and the goal is to hide the value of each X_i. Their mechanism is less general than the Wasserstein Mechanism; however it applies to a broader class of models than the Markov Quilt Mechanism. In Section 5, we do an experimental comparison between their algorithm and our method for Markov Chains and show that our method applies to a broader class of parameter values, and has utility comparable to or better than theirs.

2 The Setting

To motivate our privacy framework, we use two running examples – physical activity monitoring of a single subject and flu statistics in a social network.

Example 1: Physical Activity Monitoring. We have a time-series $X = \{X_1, X_2, \ldots\}$ where X_t denotes a discrete physical activity (e.g, running, sleeping, sitting, etc) of a subject at time t. Our

[1] An earlier version of this paper appears on Arxiv within nine days of [11]

goal is to release (an approximate) histogram of activities over a long period (say, a week), while the adversary would like to determine what the subject was doing at a time t (say, 10:30am on Jan 4).

Example 2: Flu Status. We have a dataset $X = \{X_1, \ldots, X_n\}$ where $X_i \in \{0, 1\}$ denotes person i's flu status. We also have a social network $G = (X, E)$ that describes interaction among people, and a set of probabilistic models that describe how flu spreads when people interact. Our goal is to release an approximation to $\sum_i X_i$, the number of infected people, while still ensuring privacy against an adversary who wishes to detect whether an Alice in the dataset has flu.

2.1 The Pufferfish Privacy Framework

The privacy framework of our choice is Pufferfish [19], a generalization of differential privacy [7]. A Pufferfish framework is specified by three parameters – a set $\mathcal{S}$ of secrets, a set $\mathcal{Q} \subseteq \mathcal{S} \times \mathcal{S}$ of secret pairs, and a class of data distributions Θ. $\mathcal{S}$ is the set of possible facts about the data that we might wish to hide, and could refer to a single individual's private data or part thereof. $\mathcal{Q}$ is the set of secret pairs that we wish to be indistinguishable. Finally, Θ is a set of distributions that can plausibly generate the data, and controls the amount and nature of correlation. Each $\theta \in \Theta$ represents a belief an adversary may hold about the data, and the goal of the privacy framework is to ensure indistinguishability in the face of these beliefs.

Definition 2.1. *A privacy mechanism M is said to be ϵ-Pufferfish private in a framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ if for datasets $X \sim \theta$ where $\theta \in \Theta$ and for all secret pairs $(s_i, s_j) \in \mathcal{Q}$, and for all $w \in$ Range(M), we have:*

$$e^{-\epsilon} P_{M,\theta}(M(X) = w | s_j, \theta) \leq P_{M,\theta}(M(X) = w | s_i, \theta) \leq e^{\epsilon} P_{M,\theta}(M(X) = w | s_j, \theta) \quad (1)$$

when s_i and s_j are such that $P(s_i|\theta) \neq 0$, $P(s_j|\theta) \neq 0$.

Readers familiar with differential privacy will observe that unlike differential privacy, the probability in (1) is with respect to the randomized mechanism *and* $X \sim \theta$; to emphasize this, we use the notation X instead of D. An alternative interpretation of the definition is for $X \sim \theta$, $\theta \in \Theta$, for all $(s_i, s_j) \in \mathcal{Q}$, and for all $w \in$ Range(M), we have:

$$e^{-\epsilon} \cdot \frac{P(s_i|\theta)}{P(s_j|\theta)} \leq \frac{P(s_i | M(X) = w, \theta)}{P(s_j | M(X) = w, \theta)} \leq e^{\epsilon} \cdot \frac{P(s_i|\theta)}{P(s_j|\theta)}. \quad (2)$$

In other words, knowledge of $M(X)$ does not affect the posterior ratio of the likelihood of s_i and s_j, compared to the initial belief.

[19] shows that Differential Privacy is a special case of Pufferfish, where $\mathcal{S}$ is the set of all facts of the form $s_x^i =$ *Person i has value x* for $i \in \{1, \ldots, n\}$ and x in a domain $\mathcal{X}$, $\mathcal{Q}$ is set of all pairs (s_x^i, s_z^i) for x and z in $\mathcal{X}$ with $x \neq z$, and Θ is the set of all distributions where each individual is distributed independently. Moreover, [19] shows that we cannot have both privacy and utility when Θ is the set of all distributions. As a result, it is essential to select Θ wisely; if Θ is too restrictive, then we may not have privacy against legitimate prior knowledge on the part of the adversary, and if Θ is too large, then the resulting privacy mechanisms may have little utility.

Examples. How can we model the examples we describe above in Pufferfish? To model the physical activity example, we set $\mathcal{S}$ to be the set $\{s_a^t : t = 1, \ldots, T, a \in A\}$, where A is the set of all activities and where s_a^t means that activity a happens at time t. $\mathcal{Q}$ is the set of all pairs (s_a^t, s_b^t) for a, b in A and for all t. Finally, Θ is a set of time series models that capture how people switch between activities. For example a plausible Θ is a set of Markov Chains with state space A. In the flu status example, $\mathcal{S}$ is the set $\{s_0^i, s_1^i : i = 1, \ldots, n\}$, where s_j^i means person i has disease status j (0 or 1). $\mathcal{Q}$ is $\{(s_0^i, s_1^i) : i = 1, \ldots, n\}$, which means we wish the flu status of each person to be hidden. Θ is a set of models that describe the spread of flu; a plausible Θ is a graph $G = (X, E)$ that describes personal interactions, and a set of probability distributions over the connected components of G that describes how flu spreads among interacting people.

Why Pufferfish? Perhaps the most popular framework for data privacy is differential privacy [7], where the idea is to add sufficient noise to hide the effect of participation of a single individual in the data. However, when multiple individuals' data is correlated, differential privacy is not enough to *erase the evidence of a single individual's private value* [18]. A related framework is coupled

worlds privacy [1], which accounts for correlation through a model class Θ, but still requires that the participation of a single individual (or group) in the data is indistinguishable. Finally, an extension is of differential privacy is group privacy, where the goal is to hide the effect of participation of every group of correlated individuals. We select Pufferfish over these privacy frameworks because (a) unlike differential privacy and coupled-worlds, Pufferfish is capable of hiding private values against correlation across multiple entries/individuals and (b) unlike group privacy, it is capable of providing utility in situations where a large number of entries are correlated, yet the "average amount" of correlation is small.

Concretely, for the physical activity example, both (a version of) differential privacy and coupled worlds privacy will add noise with standard deviation $\approx 1/\epsilon$ to each bin of the activity histogram. This will not erase the evidence of an activity at time t as activities are highly correlated. Since all the entries are correlated, group privacy will add noise with standard deviation $\approx T/\epsilon$ where T is the length of the chain, thus destroying all utility. In contrast, in Section 4 we show a Pufferfish mechanism which will add noise approximately equal to an upper bound on the mixing time, and thus offer both privacy and utility for rapidly mixing chains. In the flu status example, both differential privacy and coupled worlds privacy will add noise with standard deviation $\approx 1/\epsilon$ to the number of people with flu. If Alice is in a moderate sized connected component and if flu spreads in a connected component with probability $p = 0.5$, then this is not enough to hide the evidence of Alice's flu status. Group differential privacy will require addition of noise proportional to the size of the largest connected component; however, we can still achieve Pufferfish privacy by adding noise proportional to the "average spread" of flu. For a concrete numerical example, see Section 3.

3 A General Pufferfish Mechanism

There are a large number of generic privacy mechanisms [7, 21, 4, 22, 9, 3] for differential privacy. In contrast, while a number of Pufferfish mechanisms for very specific frameworks are known [19, 13], there is no completely generic mechanism that applies to *any* Pufferfish framework. Our first contribution is to provide such a mechanism. Specifically, given data represented by random variable X (that may be scalar or vector-valued), a Pufferfish privacy framework $(\mathcal{S}, \mathcal{Q}, \Theta)$, and a function F that maps X into a scalar, our goal is to design a generic mechanism M that satisfies ϵ-Pufferfish privacy in this framework and approximates $F(X)$.

Our proposed mechanism is inspired by the global sensitivity mechanism in differential privacy; recall that the latter adds noise to the function F proportional to the sensitivity, which is the worst case distance between $F(D)$ and $F(D')$ where D and D' are two datasets that differ in the value of a single individual. In Pufferfish, the quantities analogous to D and D' are the measures $P(F(X)|s_i, \theta)$ and $P(F(X)|s_j, \theta)$ for a secret pair (s_i, s_j), and therefore, the added noise should be proportional to the worst case distance between these two measures according to some metric. It turns out that the relevant metric is a form of Wasserstein distance.

Definition 3.1 (p-Wasserstein Distance)**.** *Let $(\mathcal{X}, d)$ be a Radon space, and μ, ν be two probability measures on $\mathcal{X}$ with finite p-th moment. The p-th Wasserstein distance between μ and ν is defined as:*

$$W_p(\mu, \nu) = \left(\inf_{\gamma \in \Gamma(\mu,\nu)} \int_{\mathcal{X}\times\mathcal{X}} d(x,y)^p \, \mathrm{d}\gamma(x,y) \right)^{1/p} = \left(\inf_{\gamma \in \Gamma(\mu,\nu)} \mathbb{E}[d(X,Y)^p] \right)^{1/p}, \qquad (3)$$

where $\Gamma(\mu, \nu)$ is the set of all couplings γ over μ and ν.

Intuitively, each $\gamma \in \Gamma(\mu, \nu)$ can be regarded as a way to *shift* probability mass between μ and ν; the cost of a shift γ is $(\mathbb{E}_{(X,Y)\sim\gamma}[d(X,Y)^p])^{1/p}$, and the cost of the min-cost shift is the Wasserstein distance. In particular, we are interested in the ∞-Wasserstein distance $W_\infty(\mu, \nu)$, which can be interpreted as the maximum "distance" that any probability mass moves while transforming μ to ν.

The main intuition behind our mechanism is based on this interpretation. Suppose for some (s_i, s_j) and θ, we would like to transform $P(F(X)|s_i, \theta)$ to $P(F(X)|s_j, \theta)$. Then, the maximum "distance" that any probability mass moves is $W_{i,j,\theta} = W_\infty(P(F(X)|s_i, \theta), P(F(X)|s_j, \theta))$, and adding Laplace noise with scale $W_{i,j,\theta}/\epsilon$ to F will guarantee that the likelihood ratio of the outputs under s_i and s_j lies in $[e^{-\epsilon}, e^{\epsilon}]$. Iterating over all pairs $(s_i, s_j) \in \mathcal{Q}$ and all $\theta \in \Theta$ and taking the maximum over $W_{i,j,\theta}$ leads to a mechanism for the entire framework $(\mathcal{S}, \mathcal{Q}, \Theta)$. The full mechanism is described in Algorithm 1, and its privacy properties in Theorem 3.2. Observe that when Pufferfish

reduces to differential privacy, then the corresponding Wasserstein Mechanism reduces to the global sensitivity mechanism; it is thus a generalization of the global sensitivity mechanism. A version of this mechanism for (ϵ, δ)-Pufferfish privacy is derived in Appendix A.4.

Algorithm 1 Wasserstein Mechanism (Dataset D, function F, Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$, privacy parameter ϵ)

for all $(s_i, s_j) \in \mathcal{Q}$ and any $\theta \in \Theta$ such that $P(s_i|\theta) \neq 0$ and $P(s_j|\theta) \neq 0$ **do**
  Set $\mu_{i,\theta} = P(F(X) = \cdot|s_i, \theta)$, $\mu_{j,\theta} = P(F(X) = \cdot|s_j, \theta)$. Calculate $W_\infty(\mu_{i,\theta}, \mu_{j,\theta})$
end for
Set $W = \sup_{(s_i,s_j)\in\mathcal{Q},\theta\in\Theta} W_\infty(\mu_{i,\theta}, \mu_{j,\theta})$.
return $F(D) + Z$, where $Z \sim \text{Lap}(\frac{W}{\epsilon})$

Theorem 3.2 (Privacy of the Wasserstein Mechanism)**.** *The Wasserstein Mechanism provides ϵ-Pufferfish privacy in the framework $(\mathcal{S}, \mathcal{Q}, \Theta)$.*

As a concrete example, consider the flu example with a social interaction graph of size 4. Suppose that the number N of infected individuals is described by the distribution: for any individual i, let X_i be its disease status. Then, $P(N = 0|X_i = 0) = 1/2$, $P(N = i|X_i = 0) = 1/6$ for $i \in [3]$, and moreover, $P(N = j|X_i = 1) = 1/4$ for $j \in [4]$. It is easy to show that the parameter W in Algorithm 1 is 2, and the Wasserstein Mechanism will add $\text{Lap}(2/\epsilon)$ noise to the number of infected individuals. In this case, as all the X_i's are correlated, the global sensitivity mechanism with Group Differential Privacy would add $\text{Lap}(4/\epsilon)$ noise, which gives worse utility.

Because of the extreme generality of the Pufferfish framework, it is difficult to make general statements about the utility of the Wasserstein mechanism. However, we show that the phenomenon illustrated by the preceeding example is quite general. When X can be written as $X = (X_1, \ldots, X_n)$ where X_i denotes the i-th individual's private value, and the goal is to keep each individual's value private, the Wasserstein Mechanism for Pufferfish never performs worse than the global sensitivity mechanism for the corresponding group differential privacy framework. See Appendix A.2.2 for a proof.

Theorem 3.3 (Wasserstein vs. Global Sensitivity for Group Pivacy)**.** *Let $(\mathcal{S}, \mathcal{Q}, \Theta)$ be a Pufferfish framework, and let $\mathcal{G}$ be the corresponding group differential privacy framework (so that $\mathcal{G}$ includes a group G for each set of correlated individuals in $(\mathcal{S}, \mathcal{Q}, \Theta)$). Then, for any 1-Lipschitz function F, the parameter W in the Wasserstein Mechanism for releasing F with $(\mathcal{S}, \mathcal{Q}, \Theta)$-Pufferfish privacy is less than or equal to the global sensitivity of F in the $\mathcal{G}$-group differential privacy framework.*

4 A Mechanism for Bayesian Networks

While the Wasserstein Mechanism is completely general, it may be computationally expensive in many situations. We next consider a more restricted setting where the data X can be written as a collection of variables $X = (X_1, \ldots, X_n)$ whose dependence is described by a Bayesian network, and the goal is to keep the value of each variable X_i private. We show that this setting is still practical, as it includes the physical activity monitoring example, and we conclude with a computationally efficient algorithm for this example.

The Setting. Specifically, in this section we assume that data X can be written as $X = \{X_1, \ldots, X_n\}$, where each X_i lies in a bounded domain $\mathcal{X}$. Let s_a^i denote the event that X_i takes value a. The set of secrets is $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$, and the set of secret pairs is $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [n]\}$. We also assume that there is an underlying known Bayesian network $G = (X, E)$ connecting the variables. Each $\theta \in \Theta$ that describes the distribution of the variables is based on this Bayesian network G, but may have different model parameters.

Notation. We use X with a lowercase subscript, for example, X_i, to denote a single node in G, and X with an uppercase subscript, for example, X_A, to denote a set of nodes in G. For a set of nodes X_A we use the notation $\texttt{card}(X_A)$ to denote the number of nodes in X_A. In this section, we will consider functions F that are 1-Lipschitz; that is, for all $X_1, \ldots, X_n$ and all X_i', $|F(X_1, \ldots, X_i, \ldots, X_n) - F(X_1, \ldots, X_i', \ldots, X_n)| \leq 1$. Since any Lipschitz function can be scaled to be 1-Lipschitz, this is not a significant restriction.

The Markov Quilt Mechanism. The main insight behind our mechanism is that if nodes X_i and X_j are "far apart" in the graph G, then, X_j is largely independent from X_i. Thus, to obscure the effect of X_i on the function F, it is sufficient to add noise that is proportional to the number of nodes that are "local" to X_i plus a small correction term that accounts for the effect of the distant nodes. The rest of the section will explain how to calculate this correction term, and how to determine how many nodes are local.

First, we quantify how much changing the value of a variable $X_i \in X$ can affect a set of variables $X_A \subset X$, where $X_i \notin X_A$. To this end, we define the *max-influence* of a variable X_i on a set of variables X_A as follows.

Definition 4.1. *The max-influence of a variable X_i on a set of variables X_A is defined as*

$$e(X_A|X_i) = \max_{a,b \in \mathcal{X}} \sup_{\theta \in \Theta} \max_{x_A \in \mathcal{X}^{\texttt{card}(X_A)}} \log \frac{P(X_A = x_A|X_i = a, \theta)}{P(X_A = x_A|X_i = b, \theta)}.$$

Here $\mathcal{X}$ is the range of any X_j and Θ is the set of allowable distributions. The max-influence is thus the maximum max-divergence [8] between the distributions $X_A|X_i = a, \theta$ and $X_A|X_i = b, \theta$ where the maximum is taken over any pair $(a, b) \in \mathcal{X} \times \mathcal{X}$ and any $\theta \in \Theta$. If X_A and X_i are independent, then the max-influence of X_i on X_A is 0, and a large max-influence means that changing X_i can have a large impact on the distribution of X_A.

Our mechanism will attempt to find large sets X_A such that X_i has low max-influence on X_A. The naive way to do so is through brute force search, which takes time exponential in the size of G. We next show how structural properties of the Bayesian network G can be exploited to perform this search more efficiently. For this purpose, we provide a second definition that generalizes the familiar notion of a Markov Blanket. Recall that the Markov Blanket of a node X_u in a Bayesian network consists of its parents, its children and the other parents of its children, and the rest of the nodes in the network are independent of X_u conditioned on its Markov Blanket. We define its generalization, the Markov Quilt, as follows.

Definition 4.2 (Markov Quilt)**.** *A set of nodes $X_Q, Q \subset [n]$ in a Bayesian network $G = (X, E)$ is a Markov Quilt for a node X_i if the following conditions hold:*
1. Deleting X_Q partitions G into parts X_N and X_R such that $X = X_N \cup X_Q \cup X_R$ and $X_i \in X_N$.
2. For all $x_R \in \mathcal{X}^{\texttt{card}(X_R)}$, all $x_Q \in \mathcal{X}^{\texttt{card}(X_Q)}$ and for all $a \in \mathcal{X}$, we have:

$$P(X_R = x_R|X_Q = x_Q, X_i = a) = P(X_R = x_R|X_Q = x_Q).$$

In other words, X_R is independent of X_i conditioned on X_Q.

Intuitively, X_R is a set of "remote" nodes that are far from X_i, and X_N is the set of "nearby" nodes; X_N and X_R are separated by the Markov Quilt X_Q. Observe that unlike Markov Blankets, a node can have many Markov Quilts. Figure 2 in appendix shows an example.

We can now show the following. Suppose we would like to release a 1-Lipschitz function F while protecting the value of a node X_i. If we can find a Markov Quilt (X_N, X_Q, X_R) such that the max-influence of X_i on X_Q is at most δ, then, it is sufficient to add Laplace noise to F with scale parameter $\texttt{card}(X_N)/(\epsilon - \delta)$. This motivates the following mechanism, which we call the Markov Quilt Mechanism. To protect X_i, we search over a set $S_{Q,i}$ of Markov Quilts for X_i and pick the one which requires adding the least amount of noise. We then iterate over all X_i and add to F the maximum amount of noise needed to protect any X_i; this ensures that the private values of *all nodes* are protected. Details are presented in Algorithm 2. Theorem 4.3, proved in Appendix B.3, establishes privacy of the mechanism.

Theorem 4.3 (Privacy of Markov Quilt Mechanism)**.** *If F is a 1-Lipschitz function, and if each $S_{Q,i}$ is non-empty, then the Markov Quilt Mechanism preserves ϵ-Pufferfish privacy in the framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ described in the beginning of this section.*

4.1 Case Study: Physical Activity Monitoring

Algorithm 2 can still be computationally expensive; there may be an exponential number of Markov Quilts in the graph, and enumerating over each $\theta \in \Theta$ to compute the max-influence may be computation-intensive. In this section, we show that for the physical activity monitoring example, further optimizations can be used to make the mechanism computationally efficient.

Algorithm 2 Markov Quilt Mechanism (Dataset D, function F, Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$, privacy parameter ϵ)

for all X_i **do**
  for all Markov Quilts (X_N, X_Q, X_R) for X_i in $S_{Q,i}$ **do**
    If the max-influence $e(X_Q|X_i) < \epsilon$, define the score of X_Q as: $\sigma(X_Q) = \frac{\texttt{card}(X_N)}{\epsilon - e(X_Q|X_i)}$; otherwise, $\sigma(X_Q) = \infty$.
  end for
  $\sigma_i = \min_{X_Q \in S_{Q,i}} \sigma(X_Q)$
end for
return $F(D) + (\max_i \sigma_i) \cdot Z$, where $Z \sim \text{Lap}(1)$

The Setting. Data is described by a random vector $X = \{X_1, X_2, \ldots, X_T\}$ over time $t \in [T]$ where each X_t is a state (representing an activity) that lies in a finite discrete set $\mathcal{X} = [k]$. Our goal is to hide from an adversary the value of each X_t; using s_a^t to denote the event that X_t takes value a, our secret set is $\mathcal{S} = \{s_a^t : a \in \mathcal{X}, t \in [T]\}$, and our set of secret pairs is $\mathcal{Q} = \{(s_a^t, s_b^t) : a, b \in \mathcal{X}, a \neq b, t \in [T]\}$. The underlying Bayesian network G is a Markov chain $X_1 \to X_2 \to \ldots \to X_T$. Each $\theta \in \Theta$ corresponds to a $k \times k$ transition matrix P_θ for this Markov chain.

Efficient Search over Markov Quilts. Searching over Markov Quilts is efficient in this case; as we show in Lemma B.2 in the Appendix, the minimal Markov Quilts for a node X_i in a Markov Chain comprise of:

$$S_{Q,i} = \{\{X_{i-a}, X_{i+b}\}, \{X_{i-a}\}, \{X_{i+b}\}, \emptyset, 1 \leq a \leq i-1, 1 \leq b \leq T-i\}$$

Thus, we need to search over $\mathcal{O}\left(T^2\right)$ Markov Quilts for each X_i.

Efficient Search over Θ. Even with a limited number of Markov Quilts, the mechanism may be computationally intractible if we need to search over a large space of $\theta \in \Theta$. This may happen when the number of states is large.

We next show that under some assumptions, this search is also considerably simplified in Markov Chains. For any $\theta \in \Theta$, let π_θ be the stationary distribution of the Markov Chain whose transition matrix is P_θ. We define: $\pi_\Theta = \min_{x \in \mathcal{X}, \theta \in \Theta} \pi_\theta(x)$ as the probability of the least probable state in the stationary distribution corresponding to any $\theta \in \Theta$. Additionally, we define: $g_\Theta = \min_{\theta \in \Theta} \min\{1 - |\lambda| : P_\theta x = \lambda x, \lambda < 1\}$ as the least eigengap of any P_θ.

Suppose the Markov Chains induced by each $P_\theta \in \Theta$ are irreducible, aperiodic and reversible with $\pi_\Theta > 0$ and $g_\Theta > 0$. Then, we can provide upper bounds on the max-influence of a Markov Quilt $X_Q = \{X_{i-a}, X_{i+b}\}$ on a node X_i based on Markov Chain theory:

$$e(X_Q|X_i) \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right) + 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).$$

Similar results may be obtained for Markov Quilts of the form $X_Q = \{X_{i-a}\}$ or $\{X_{i+b}\}$ as well. Details are in Lemma B.4 in the Appendix.

When Θ is parametrized by g_Θ and π_Θ, (an upper bound on) the score of each X_Q may thus be calculated in $\mathcal{O}(1)$ time based on Lemma B.4. Thus, the computational complexity of the Markov Quilt mechanism is $\mathcal{O}(T^3)$. A more optimized version which runs in $\mathcal{O}(T^2)$ time in some good cases is derived in Appendix B.8.2.

Utility. We conclude this section with an analysis of the utility of Algorithm 2 for this case study.

Theorem 4.4 (Utility of Algorithm 2 for Physical Activity Monitoring)**.** *Suppose we apply Algorithm 2 to release an approximation to a* 1*-Lipschitz function F of the states in the physical activity monitoring example. If π_Θ and g_Θ are as defined above, and if the length T of the chain satisfies $T \geq 2\lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil + 3$, then the Laplace noise added by the Markov Quilt Mechanism has scale parameter $\leq C/\epsilon$ for some positive constant C that depends on Θ only.*

The proof is in the Appendix. Observe that unlike group privacy, the amount of noise added does not grow with T; a careful look at the proof shows that the amount of noise added is approximately

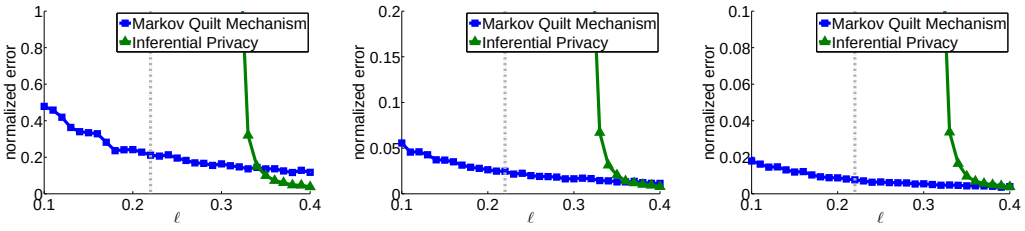

(a) $\epsilon = 1$. (b) $\epsilon = 5$. (c) $\epsilon = 10$.

Figure 1: Normalized difference between private and non-private sum (averaged over 500 runs) vs. minimum value of the interval ℓ. The y-axis is scaled to $1/\epsilon$ for better comparison. The inferential privacy method does not work in the region left to the grey dashed line.

an upper bound on the mixing time of the chain. Thus if Θ consists of rapidly mixing chains, then Algorithm 2 provides both privacy and utility.

5 Simulations

We next illustrate the performance of the Markov Quilt Mechanism by applying it to a simple example and comparing it with the inferential privacy algorithm provided by concurrent work [11]. This work provides an alternative algorithm for Pufferfish privacy when the data X can be written as $X = (X_1, \ldots, X_n)$ and the goal is to protect the value of each individual X_i.

Specifically, we consider the setting in Section 4.1 where the underlying Bayesian network is a discrete time homogeneous Markov chain with two states $\{0, 1\}$. Observe that any such chain may be described by two parameters – $p_0 = P(X_{i+1} = 0 | X_i = 0)$ and $p_1 = P(X_{i+1} = 1 | X_i = 1)$ – which completely determine the transition matrix of the chain. A distribution θ is thus represented by a tuple (p_0, p_1) and a distribution family Θ by an interval $\Theta = [\ell, u]$, which includes any $p_0, p_1 \in [\ell, u]$. Observe that when $0 < p_0, p_1 < 1$, the chain is guaranteed to be aperiodic, irreducible and reversible. The function F is the summation S of states over time $t = 1, \ldots, T$, which is 1-Lipschitz.

To generate synthetic data from a distribution family $\Theta = [\ell, u]$, we pick p_0 and p_1 uniformly from $[\ell, u]$, and an initial state distribution uniformly from the probability simplex. We then generate a state sequence $X_1, \ldots, X_T$ of length T from the corresponding distribution. Both the Inferential Privacy method and the Markov Quilt Mechanism are run on the resulting chain; for each method, we compute the amount of noise needed for each $\theta \in \Theta$ and then add the maximum amount of noise needed in the process. This can be done through grid search as the parameter space is small.

Utility is measured by the difference between S, the sum $\sum_t X_t$ and S', the private version of this sum, normalized by the chain length T; a smaller difference implies higher utility. In our simulations, we fix the length of the chain $T = 100$, and investigate the effect of Θ and ϵ on utility. For ease of presentation, we restrict Θ to be intervals where $u = 1 - \ell$, and vary ℓ from 0.1 to 0.4. We vary ϵ in $\{1, 5, 10\}$, repeat each experiment 500 times and report the averaged value of $|S' - S|/T$.

Figure 1 shows the results. For any given ϵ, the error of both algorithms decrease as ℓ increases, i. e, the distribution class Θ becomes more narrow. When ℓ is left to the grey dashed line, the algorithm of [11] does not work (note that this algorithm only works when the spectral norm of a certain influence matrix is less than 1); however the Markov Quilt Mechanism still provides privacy and reasonable utility. The results are mixed in the region where [11] does work; for ℓ close to 0.5 and for small ϵ, [11] performs a little better, whereas for smaller ℓ, the Markov Quilt Mechanism is significantly better. Thus, we conclude that for this problem, the Markov Quilt Mechanism applies to a much wider range of distribution families; in the region where both mechanisms work, the Markov Quilt Mechanism performs significantly better than the method of [11] for a certain range of parameter values, and slightly worse for the remaining range.

References

[1] Raef Bassily, Alex Groce, Justin Katz, and Adam Smith. Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 439–448. IEEE, 2013.

[2] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization, revisited. arXiv:1405.7085, 2014.

[3] Kamalika Chaudhuri, Daniel Hsu, and Shuang Song. The large margin mechanism for differentially private maximization. In *NIPS*, 2014.

[4] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

[5] Kamalika Chaudhuri, Anand D. Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *Advances in Neural Information Processing Systems*, pages 998–1006, 2012.

[6] John Duchi, Martin J. Wainwright, and Michael Jordan. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances in Neural Information Processing Systems 26*, pages 1529–1537, 2013.

[7] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, 2006.

[8] C. Dwork, G. Rothblum, and S. Vadhan. Boosting and differential privacy. In *FOCS*, 2010.

[9] Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.

[10] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.

[11] Arpita Ghosh and Robert Kleinberg. Inferential privacy guarantees for differentially private mechanisms. *arXiv preprint arXiv:1603.01508*, 2016.

[12] M. Hardt and A. Roth. Beyond worst-case analysis in private singular vector computation. In *STOC*, 2013.

[13] Xi He, Ashwin Machanavajjhala, and Bolin Ding. Blowfish privacy: tuning privacy-utility trade-offs using policies. In *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*, pages 1447–1458, 2014.

[14] Hafiz Imtiaz and Anand D. Sarwate. Symmetric matrix perturbation for differentially private pca. In *ICASSP*, 2015.

[15] P. Jain, P. Kothari, and A. Thakurta. Differentially private online learning. In *COLT*, 2012.

[16] Stephan Kessler, Erik Buchmann, and Klemens Böhm. Deploying and evaluating pufferfish privacy for smart meter data. *Karlsruhe Reports in Informatics*, 1, 2015.

[17] D. Kifer, A. Smith, and A. Thakurta. Private convex optimization for empirical risk minimization with applications to high-dimensional regression. In *COLT*, 2012.

[18] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 193–204, 2011.

[19] Daniel Kifer and Ashwin Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Trans. Database Syst.*, 39(1):3, 2014.

[20] David Levin, Yuval Peres, and Elizabeth Wilmer. *Markov Chains and Mixing Time*. American Mathematical Society, 2009.

[21] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *FOCS*, 2007.

[22] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing (STOC '07)*, pages 75–84, New York, NY, USA, 2007. ACM.

[23] A.D. Sarwate and K. Chaudhuri. Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data. *Signal Processing Magazine, IEEE*, 30(5):86–94, Sept 2013.

[24] Yu-Xiang Wang, Stephen E. Fienberg, and Alexander J. Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 2493–2502, 2015.

[25] Bin Yang, Issei Sato, and Hiroshi Nakagawa. Bayesian differential privacy on correlated data. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 747–762. ACM, 2015.

A Details for Wasserstein Mechanism

A.1 Proof to Theorem 3.2 (Privacy of the Wasserstein Mechanism)

Proof. Let (s_i, s_j) be a pair of secrets in $\mathcal{Q}$ and let $\theta \in \Theta$ such that $P(s_i|\theta) > 0$ and $P(s_j|\theta) > 0$. Recall that $\mu_{i,\theta} = P(F(X) = \cdot|s_i, \theta)$, and $\mu_{j,\theta}$ is defined similarly. Let $\gamma^* = \gamma^*(\mu_{i,\theta}, \mu_{j,\theta})$ be the coupling between $\mu_{i,\theta}$ and $\mu_{j,\theta}$ that achieves the ∞-Wasserstein distance.

Using M to denote the Wasserstein mechanism, we can write for any w,

$$\begin{aligned}
&\frac{P(M(X) = w|s_i, \theta)}{P(M(X) = w|s_j, \theta)} \\
= \;& \frac{\int_t P(F(X) = t|s_i, \theta)P(Z = w - t)dt}{\int_s P(F(X) = s|s_j, \theta)P(Z = w - s)ds} \\
= \;& \frac{\int_t P(F(X) = t|s_i, \theta)e^{-\epsilon|w-t|/W}dt}{\int_s P(F(X) = s|s_j, \theta)e^{-\epsilon|w-s|/W}ds} \\
= \;& \frac{\int_t \int_{s=t-W}^{t+W} \gamma^*(t, s)e^{-\epsilon|w-t|/W}dsdt}{\int_s \int_{t=s-W}^{s+W} \gamma^*(t, s)e^{-\epsilon|w-s|/W}dtds}. \qquad (4)
\end{aligned}$$

Here the first step follows from the definition of the Wasserstein mechanism, the second step from properties of the Laplace distribution, and the third step from the following property of ∞-Wasserstein distance:

$$\begin{aligned}
P(F(X) = t|s_i, \theta) &= \int_{s=t-W}^{t+W} \gamma^*(t, s)ds, \\
P(F(X) = s|s_j, \theta) &= \int_{t=s-W}^{s+W} \gamma^*(t, s)dt.
\end{aligned}$$

Observe that in the last step of (4), $|s - t| \leq W$ in both the numerator and denominator integrals; therefore we have the following inequality

$$\frac{\int_t \int_{s=t-W}^{t+W} \gamma^*(t, s)e^{-\epsilon|w-t|/W}dsdt}{\int_s \int_{t=s-W}^{s+W} \gamma^*(t, s)e^{-\epsilon|w-s|/W}dtds} \leq e^{\epsilon} \frac{\int_t \int_{s=t-W}^{t+W} \gamma^*(t, s)e^{-\epsilon|w-s|/W}dsdt}{\int_s \int_{t=s-W}^{s+W} \gamma^*(t, s)e^{-\epsilon|w-s|/W}dtds}.$$

By definition of ∞-Wasserstein distance, $\gamma^*(t, s) = 0$ when $|s - t| > W$. Therefore, the right hand side of the above inequality is equal to

$$e^{\epsilon} \frac{\int_t \int_s \gamma^*(t, s)e^{-\epsilon|w-s|/W}dsdt}{\int_s \int_t \gamma^*(t, s)e^{-\epsilon|w-s|/W}dtds} \leq e^{\epsilon}.$$

Similar argument can be used to show that $\frac{P(M(X)=w|s_i,\theta)}{P(M(X)=w|s_j,\theta)} \geq e^{-\epsilon}$. Combining these two facts concludes the proof of the theorem. □

A.2 Comparison of Wasserstein Mechanism with Global Sensitivity Mechanism for Group DP

In the flu spread example in Section 3, Wasserstein Mechanism with Pufferfish privacy adds noise with smaller variance compared to what Global Sensitivity Mechanism with Group DP does. The intuition behind such difference is that Wasserstein Mechanism with Pufferfish privacy exploits information on correlation and thus add small noise when correlation is low, while Global Sensitivity Mechanism with Group DP treats individuals only as either correlated or uncorrelated. In this section, we show that Wasserstein Mechanism always yields utility better than or equal to its corresponding Global Sensitivity counterpart.

A.2.1 Overview of Group DP Framework and Mechanisms

A group DP framework $\mathcal{G}$ considers a set of groups $\{G_1, \cdots, G_m\}$, where each group G_k is a set of individuals, and the goal is to hide the effect of participation of each group G_k. For each G_k, we define a set $X_k \subset X$ representing private values in X contributed by individuals in G_k. Let D_k be set of input pairs (x, y) such that x and y differ only by entries in X_k. Then

Definition A.1 (Group DP)**.** *Under group DP framework $\mathcal{G}$, a randomized algorithm M is called ϵ-group differentially private if for all group G_k, all $(x, y) \in D_k$ and all $w \in Range(M)$,*

$$P[M(x) = w] \leq \exp(\epsilon) P[M(y) = w].$$

Global Sensitivity Mechanism can be used to guarantee group DP, with the sensitivity now defined with respect to groups in $\mathcal{G}$. Let F be a function that maps input random variables to a number.

Definition A.2 (Global Sensitivity of F to a group G_k)**.** *Global sensitivity of F to changes in G_k is*

$$\Delta_{G_k} F = \max_{(x,y) \in D_k} |F(x) - F(y)|,$$

that is, the maximum change in F's value when entries in X_k change.

Then we can define the Global Sensitivity of a Group DP framework $\mathcal{G}$.

Definition A.3 (Global Sensitivity of F in $\mathcal{G}$)**.** *Global Sensitivity of F in a Group DP framework $\mathcal{G}$ is*

$$\Delta_{\mathcal{G}} F = \max_{k \in [m]} \Delta_{G_k} F.$$

Definition A.4 (Global Sensitivity Mechanism for Group DP)**.** *Given a function F, a group DP framework $\mathcal{G}$, an input x and a privacy parameter ϵ, the Global Sensitivity Mechanism is defined as*

$$M(x, F(\cdot), \epsilon, \mathcal{G}) = F(x) + Z$$

where $Z \sim Lap(\Delta_{\mathcal{G}} F / \epsilon)$.

Theorem A.5. *Global Sensitivity Mechanism $M(F(\cdot), x, \epsilon, \mathcal{G})$ guarantees ϵ-group DP for function F under framework $\mathcal{G}$.*

Proof. Let w be any number in M's range, x, y be two inputs differed by entries from individuals within in a group G_k, then we have

$$\begin{aligned}
\frac{p(M(x) = w)}{p(M(y) = w)} &= \frac{p(Z = w - F(x))}{p(Z = w - F(y))} \\
&= \frac{\exp(-\frac{\epsilon |F(x) - w|}{\Delta_{\mathcal{G}} F})}{\exp(-\frac{\epsilon |F(y) - w|}{\Delta_{\mathcal{G}} F})} \\
&= \exp\left(\frac{\epsilon(|F(y) - w| - |F(x) - w|)}{\Delta_{\mathcal{G}} F}\right) \\
&\leq \exp(\frac{\epsilon |F(y) - F(x)|}{\Delta_{\mathcal{G}} F}) \\
&\leq \exp(\epsilon).
\end{aligned}$$

The first inequality follows from the triangle inequality and the second from definition of global sensitivity. □

A.2.2 Proof to Theorem 3.3

Before proving the Theorem, we formally define the correspondence between Pufferfish framework and Group DP framework.

Consider a Pufferfish framework that keeps each individuals value private from an adversary. Let $[n]$ be the set of all individuals, i be the i-th individual, and $X = \{X_1, \cdots, X_n\}$ be the set of input random variables with each X_i being a single individual i's private value. Let $\mathcal{X}$ be the domain of X_is in X. Define the set of secrets as $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$ where s_a^i is the event that X_i takes

value a, and define the set of secret pairs as $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [n]\}$. Let Θ be the set of possible distributions of X. The Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ with parameters defined above protects the privacy of each individual's value X_i for any X following any distribution $\theta \in \Theta$.

Now we formally define the corresponding Group DP framework $\mathcal{G}$ that contains several groups, each representing a set of correlated individuals in Θ. The corresponding Group DP framework $\mathcal{G}$ considers a partition $G_1, \cdots, G_m$ of all individuals $[n]$ such that private values of individuals from different groups are independent. Formally, for any two groups G_p, G_q with $p \neq q$, and any $i \in G_p, j \in G_q$, the value of X_i and X_j are independent. For the Global Sensitivity Mechanism to get the best utility, each group G_k should be "minimal" in the sense that the value of any individual in G_k must be correlated with that of some other individual in G_k. One can imagine an undirected graph in which each node represents an individual and an edge exists between two individuals if their private values are correlated. Groups in the partition $\{G_1, \cdots, G_m\}$ corresponds to connected components in that graph. Global Sensitivity Mechanism for Group DP will add noise drawn from $\text{Lap}(\Delta_{\mathcal{G}} F/\epsilon)$.

Proof. To guarantee ϵ-Pufferfish privacy under $(\mathcal{S}, \mathcal{Q}, \Theta)$, Wasserstein Mechanism for Pufferfish framework will add noise drawn from $\text{Lap}(W/\epsilon)$. The parameter W in Wasserstein Mechanism (Algorithm 1) is

$$W = \sup_{(s_a^i, s_b^i) \in \mathcal{Q}, \theta \in \Theta} W_\infty(\mu_{s_a^i, \theta}, \mu_{s_b^i, \theta}).$$

where $\mu_{i,a,\theta} = P(F(X) = \cdot | X_i = a, \theta)$ and $\mu_{i,a,\theta} = P(F(X) = \cdot | X_i = a, \theta)$.

In the corresponding Group DP framework $\mathcal{G}$ described above, Global Sensitivity Mechanism adds noise drawn from $\text{Lap}(\Delta_{\mathcal{G}} F/\epsilon)$. Now we will prove that $W \leq \Delta_{\mathcal{G}} F$.

For any secret pair $(s_a^i, s_b^i) \in \mathcal{Q}$, we can find the group G_k such that $i \in G_k$. Let X_S be the set of private values associated with individuals in G_k, and X_T be the set of the other private values in X. We observe that $X_S \bigcup X_T = X$ and $X_S \bigcap X_T = \emptyset$. And since values in X_T does not depend on X_i, we have $P(X_T = x_T | X_i = a, \theta) = P(X_T = x_T | X_i = b, \theta) = P(X_T = x_T | \theta)$ for all realization x_T of X_T.

Let

$$\begin{aligned}\nu_{i,a,\theta,x_T} =& P(F(X_S, X_T = x_T) = \cdot | X_i = a, \theta) \\ \nu_{i,b,\theta,x_T} =& P(F(X_S, X_T = x_T) = \cdot | X_i = b, \theta).\end{aligned}$$

Marginalize over $p(X_T = x_T | \theta)$, we have

$$\begin{aligned}\mu_{i,a,\theta} &= \sum_{x_T} P(X_T = x_T | \theta) \nu_{i,a,\theta,x_T}, \\ \mu_{i,b,\theta} &= \sum_{x_T} P(X_T = x_T | \theta) \nu_{i,b,\theta,x_T}.\end{aligned}$$

Notice that this two probability distributions are mixtures of ν_{i,a,θ,x_T}s and ν_{i,b,θ,x_T}s receptively, with the same set of mixing coefficients $p(X_T = x_T | \theta)$s. By Lemma A.6, we have

$$W_\infty(\mu_{i,a,\theta}, \mu_{i,b,\theta}) \leq \max_{x_T} W_\infty(\nu_{i,a,\theta,x_T}, \nu_{i,b,\theta,x_T}).$$

By definition, ∞-Wasserstein of two distributions can be upper bounded by the range of the union of the two supports, i.e.,

$$\begin{aligned}W_\infty(\nu_{i,a,\theta,x_T}, \nu_{i,b,\theta,x_T}) \leq& \max_{x_S, x_S'} |F(X_S = x_S, X_T = x_T | s_a^i, \theta) - F(X_S = x_S', X_T = x_T | s_b^i, \theta)| \\ \leq& \Delta_{G_k} F,\end{aligned}$$

where the last inequality follows from definition of global sensitivity with respect to G_k. Therefore we have

$$W_\infty(\mu_{i,a,\theta}, \mu_{i,b,\theta}) \leq \Delta_{|\mathcal{G}|} F.$$

Since this holds for all $(s_a^i, s_a^i) \in Q$ and all $\theta \in \Theta$, we have

$$W \leq \Delta_{|\mathcal{G}|} F.$$

□

Lemma A.6. *Let $\{\mu_i\}_{i=1}^n$, $\{\nu_i\}_{i=1}^n$ be two sets of probability distributions. Let $\mu = \sum_i c_i\mu_i$ and $\nu = \sum_i c_i\nu_i$ be two probability mixtures with shared mixing coefficients c_is such that $c_i \geq 0, \sum_{i=1}^n c_i = 1$. Then $W_\infty(\mu, \nu) \leq \max_{i\in[n]} W_\infty(\mu_i, \nu_i)$.*

Proof. For all $i \in [n]$, let γ_i be the coupling between μ_i and ν_i that achieves $W_\infty(\mu_i, \nu_i)$. Let the support of μ, ν be S, and that of μ_i, ν_i be S_i. We can extend each coupling γ_i to have value 0 at $(x, y) \in S\backslash S_i$.

We can construct a coupling γ between μ and ν as follows:

$$\gamma(x, y) = \sum_{i=1}^{n} c_i\gamma_i(x, y).$$

Since this is a valid coupling between μ and ν, $W_\infty(\mu, \nu)$ must be no larger than the distance induced by the specific coupling, which is $\max_i W_\infty(\mu_i, \nu_i)$. Therefore we have

$$W_\infty(\mu, \nu) \leq \max_{i\in[n]} W_\infty(\mu_i, \nu_i).$$

□

A.3 Optimality Example for Wasserstein Mechanism

Suppose we are given a function F and Pufferfish parameters $(\mathcal{S}, \mathcal{Q}, \Theta)$ and there exists $(s_i, s_j) \in \mathcal{Q}$ and $\theta \in \Theta$, such that for some positive value Δ, we have

$$P(F(X) = t|s_i, \theta) = \begin{cases} 1 & , t = 0 \\ 0 & , \text{otherwise} \end{cases}$$

and

$$P(F(X) = t|s_j, \theta) = \begin{cases} 1 - \delta & , t = 0 \\ \delta & , t = \Delta \\ 0 & , \text{otherwise.} \end{cases}$$

Suppose we add noise $Z \sim \text{Lap}(\alpha)$. Then for any w, any secret pair $(s_i, s_j) \in \mathcal{Q}$ and any distribution $\theta \in \Theta$, we want $e^{-\epsilon} \leq \frac{P(F(X)+Z=w|s_j,\theta)}{P(F(X)+Z=w|s_i,\theta)} \leq e^\epsilon$. Taking $w = \Delta$, and the specific secret pair (s_i, s_j) and distribution θ mentioned above, we have

$$\begin{aligned} &\frac{P(F(X) + Z = \Delta|s_j, \theta)}{P(F(X) + Z = \Delta|s_i, \theta)} \\ =&\frac{(1 - \delta)P(Z = \Delta) + \delta P(Z = 0)}{P(Z = \Delta)} \\ =&\frac{(1 - \delta)e^{-\Delta/\alpha} + \delta}{e^{-\Delta/\alpha}} \\ =&(1 - \delta) + \delta e^{\Delta/\alpha}. \end{aligned}$$

For large Δ, this value is of order $\mathcal{O}\left(e^{\Delta/\alpha}\right)$. So to guarantee ϵ-Pufferfish privacy, we need to have $\alpha = \Omega\left(\Delta/\epsilon\right)$.

A.4 Wasserstein Mechanism with (ϵ, δ)-Pufferfish Privacy

Definition A.7. *A privacy mechanism M is said to be (ϵ, δ)-Pufferfish private with Pufferfish parameters $(\mathcal{S}, \mathcal{Q}, \Theta)$ if for datasets $X \sim \theta$ where $\theta \in \Theta$ and for all secret pairs $(s_i, s_j) \in \mathcal{Q}$, and for all $S \subseteq$ Range(M), we have:*

$$\begin{aligned} P(M(X) \in S|s_i, \theta) &\leq e^\epsilon P(M(X) \in S|s_j, \theta) + \delta \\ P(M(X) \in S|s_j, \theta) &\leq e^\epsilon P(M(X) \in S|s_i, \theta) + \delta \end{aligned} \tag{5}$$

when s_i and s_j are such that $P(s_i|\theta) \neq 0$, $P(s_j|\theta) \neq 0$.

We provide a variant of Wasserstein Mechanism which guarantees (ϵ, δ)-Pufferfish Privacy. This mechanism has significantly improved utility compared to the original Wasserstein Mechanism when the distribution of $F(X)$ is extremely light tailed, and the Wasserstein distance is dominated by the coupling in the tail part.
Consider the following scenario. The distribution of $F(X)$ is continuous and has support on $[0, 100]$. There exists a secret pair (s_i, s_j), such that for some small $\delta \in (0, 1)$,

$$P(F(X) = t|s_i, \theta) = \begin{cases} 1-\delta & ,t \in [0,1] \\ \delta/99 & ,t \in (1,100] \end{cases}$$

and

$$P(F(X) = t|s_j, \theta) = \begin{cases} 1-\delta & ,t \in [0,1] \\ \delta & ,t \in (1,2] \\ 0 & ,t \in (2,100] \end{cases}.$$

It is obvious that the ∞-Wasserstein distance between $F(X) = t|s_i, \theta$ and $F(X) = t|s_j, \theta$ is at least 98. However, if we are allowed to "ignore" the tails of the two distributions which have probability mass δ, we get two identical "distributions" with Wasserstein distance equal to 0. We can thus hope for better utility under (ϵ, δ)-Pufferfish privacy for such case.

We build an algorithm, called (ϵ, δ)-Wasserstein Mechanism, that guarantees (ϵ, δ)-Pufferfish privacy.

1. **Inputs:** Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$, privacy parameter (ϵ, δ), function F (with continuous cdf), dataset D.
2. For any $(s_i, s_j) \in \mathcal{Q}$ and any $\theta \in \Theta$ such that $P(s_i|\theta) \neq 0$ and $P(s_j|\theta) \neq 0$, find $T_i, T_j \subseteq \mathbb{R}$ such that $P(F(X) \in T_i|s_i, \theta) = \delta$, $P(F(X) \in T_j|s_j, \theta) = \delta$. Let

$$\nu_{i,\theta,T_i}(t) = \begin{cases} P(F(X) = t|s_i, \theta)/(1-\delta) & ,t \notin T_i \\ 0 & ,t \in T_i \end{cases}$$

   and

$$\nu_{j,\theta,T_j}(t) = \begin{cases} P(F(X) = t|s_j, \theta)/(1-\delta) & ,t \notin T_j \\ 0 & ,t \in T_j. \end{cases}$$

3. Let $W = \sup_{(s_i,s_j)\in\mathcal{Q},\theta\in\Theta} \inf_{T_i,T_j} W_\infty(\nu_{i,\theta,T_i}, \nu_{j,\theta,T_j})$.
4. Output $F(D) + Z$, where $Z \sim \text{Lap}(\frac{W}{\epsilon})$.

Theorem A.8. *The (ϵ, δ)-Wasserstein Mechanism guarantees (ϵ, δ)-Pufferfish privacy in the framework $(\mathcal{S}, \mathcal{Q}, \Theta)$.*

Proof. Let (s_i, s_j) be any pair of secrets in $\mathcal{Q}$ and let $\theta \in \Theta$ such that $P(s_i|\theta) > 0$ and $P(s_j|\theta) > 0$. Since $F(X)$ has a continuous cumulative distribution function, for any $\delta \in [0, 1]$, we can always find $T_i, T_j \subseteq \mathbb{R}$ such that $P(F(X) \in T_i|s_i, \theta) = \delta$, $P(F(X) \in T_j|s_j, \theta) = \delta$. Denote $\bar{T}_i = \mathbb{R}\backslash T_i, \bar{T}_j = \mathbb{R}\backslash T_j$.
For a set $S \subseteq \mathbb{R}$ and a value $t \in \mathbb{R}$, denote $S - t = \{w - t|w \in S\}$.
Let M denote the (ϵ, δ)-Wasserstein mechanism. For any $S \subseteq \mathbb{R}$ and any T_i, T_j that satisfies the above condition, we have

$$\begin{aligned}
&\frac{P(M(X) \in S|s_i, \theta) - \delta}{P(M(X) \in S|s_j, \theta)} \\
=&\frac{\int_t P(F(X) = t|s_i, \theta)P(Z \in S - t)dt - \delta}{\int_t P(F(X) = t|s_j, \theta)P(Z \in S - t)dt} \\
=&\frac{\int_{t\in T_i} P(F(X) = t|s_i, \theta)P(Z \in S - t)dt + \int_{t\in\bar{T}_i} P(F(X) = t|s_i, \theta)P(Z \in S - t)dt - \delta}{\int_{t\in T_j} P(F(X) = t|s_j, \theta)P(Z \in S - t)dt + \int_{t\in\bar{T}_j} P(F(X) = t|s_j, \theta)P(Z \in S - t)dt} \\
\leq&\frac{\int_{t\in\bar{T}_i} P(F(X) = t|s_i, \theta)P(Z \in S - t)dt}{\int_{t\in\bar{T}_j} P(F(X) = t|s_j, \theta)P(Z \in S - t)dt},
\end{aligned} \tag{6}$$

where the inequality follows from

$$\int_{t\in T_i} P(F(X)=t|s_i,\theta)P(Z\in S-t)dt \le \int_{t\in T_i} P(F(X)=t|s_i,\theta)dt = \delta,$$

and

$$\int_{t\in T_j} P(F(X)=t|s_j,\theta)P(Z\in S-t)dt \ge 0.$$

Recall that $\nu_{i,\theta,T_i}(t) = p(F(X)=t|s_i,\theta)/(1-\delta)$ for $t\in\bar{T}_i$ and $\nu_{i,\theta,T_i}(t)=0$ otherwise, and $\nu_{j,\theta,T_j}(t) = P(F(X)=t|s_j,\theta)/(1-\delta)$ for $t\in\bar{T}_j$ and $\nu_{j,\theta,T_j}(t)=0$ otherwise. Notice that the normalization factor $1-\delta$ is simply used to make ν_{i,θ,T_i} and ν_{j,θ,T_j} distribution functions.

Then for any $w\in\mathbb{R}$, we have

$$\begin{aligned}
&\frac{\int_{t\in\bar{T}_i} P(F(X)=t|s_i,\theta)P(Z=w-t)dt}{\int_{t\in\bar{T}_j} P(F(X)=t|s_j,\theta)P(Z=w-t)dt}\\
=&\frac{\int_t (1-\delta)\nu_{i,\theta,T_i}(t)P(Z=w-t)dt}{\int_s (1-\delta)\nu_{j,\theta,T_j}(s)P(Z=w-s)ds}\\
=&\frac{\int_t \nu_{i,\theta,T_i}(t)P(Z=w-t)dt}{\int_s \nu_{j,\theta,T_j}(s)P(Z=w-s)ds}\\
\le& e^{\epsilon},
\end{aligned}$$

where the last step follows from the exact same analysis used in the proof of Theorem 3.2 in Section A.1.
Since (6) can be written as

$$\begin{aligned}
&\frac{\int_{t\in\bar{T}_i} P(F(X)=t|s_i,\theta)\int_{w\in S} P(Z=w-t)dwdt}{\int_{t\in\bar{T}_j} P(F(X)=t|s_j,\theta)\int_{w\in S} P(Z=w-t)dwdt}\\
=&\frac{\int_{w\in S}\int_{t\in\bar{T}_i} P(F(X)=t|s_i,\theta)P(Z=w-t)dtdw}{\int_{w\in S}\int_{t\in\bar{T}_j} P(F(X)=t|s_j,\theta)P(Z=w-t)dtdw},
\end{aligned}$$

which is upper bounded by e^{ϵ}, we can conclude that for any S, any (s_i,s_j) pairs and any T_i,T_j satisfying $P(F(X)\in T_i|s_i,\theta)=\delta$, $P(F(X)\in T_j|s_j,\theta)=\delta$,

$$\frac{P(M(X)\in S|s_i,\theta)-\delta}{P(M(X)\in S|s_j,\theta)} \le e^{\epsilon}.$$

The same analysis hold when we interchange s_i and s_j.

□

B Details for Markov Quilt Mechanism

B.1 Illustration of Markov Quilt

Figure 2 illustrate the difference between Markov Blanket and Markov Quilt as defined in Definition 4.2.

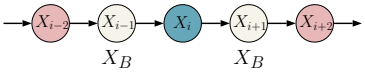

(a) The Markov Blanket of X_i is $X_B = \{X_{i-1}, X_{i+1}\}$

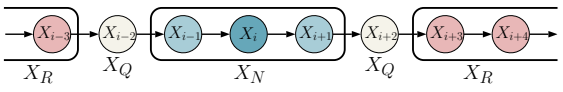

(b) One Markov Quilt of X_i is $X_Q = \{X_{i-2}, X_{i+2}\}$ with corresponding $X_N = \{X_{i-1}, X_i, X_{i+1}\}$, $X_R = \{\dots, X_{i-3}, X_{i+3}, \dots\}$

Figure 2: Illustration of Markov Blanket and Markov Quilt

B.2 Generalization to Vector-valued Functions

The Markov Quilt Mechanism can be easily generalized to vector-valued functions F. If F is 1-Lipschitz with respect to L_1 norm, then from Proposition 1 of [7], adding noise drawn from $\sigma_{\max} \cdot \text{Lap}(1)$ to each output dimension of F guarantees ϵ-Pufferfish privacy, where $\sigma_{\max}$ is the quantity in the Markov Quilt Mechanism.

B.3 Proof to Theorem 4.3 (Privacy of the Markov Quilt Mechanism)

Proof. Pick a specific secret pair $(s_a^i, s_b^i) \in \mathcal{Q}$ and a fixed $\theta \in \Theta$. Let X_Q be the Markov Quilt for X_i which has the minimum score $\sigma(X_Q)$, and suppose that deleting X_Q breaks up the underlying Bayesian network into X_N and X_R where $X_i \in X_N$.

For any w we have

$$\begin{aligned}
&\frac{P(F(X) + \sigma_{\max} \cdot Z = w | X_i = a, \theta)}{P(F(X) + \sigma_{\max} \cdot Z = w | X_i = b, \theta)} \\
=&\frac{\int_{x_{R\cup Q}} P(F(X) + \sigma_{\max} Z = w | X_i = a, X_{R\cup Q} = x_{R\cup Q}, \theta) P(X_{R\cup Q} = x_{R\cup Q} | X_i = a, \theta) dx_{R\cup Q}}{\int_{x_{R\cup Q}} P(F(X) + \sigma_{\max} Z = w | X_i = b, X_{R\cup Q} = x_{R\cup Q}, \theta) P(X_{R\cup Q} = x_{R\cup Q} | X_i = b, \theta) dx_{R\cup Q}}.
\end{aligned} \tag{7}$$

Recall that for any value $x_{R\cup Q}$ of $X_{R\cup Q}$,

$$\frac{P(X_{R\cup Q} = x_{R\cup Q} | X_i = a, \theta)}{P(X_{R\cup Q} = x_{R\cup Q} | X_i = b, \theta)} = \frac{P(X_R = x_R | X_Q = x_Q, X_i = a, \theta) P(X_Q = x_Q | X_i = a, \theta)}{P(X_R = x_R | X_Q = x_Q, X_i = b, \theta) P(X_Q = x_Q | X_i = b, \theta)}.$$

Since X_Q is a Markov Quilt for X_i and $X_i \notin X_R$, we have $P(X_R | X_Q, X_i = a, \theta) = P(X_R | X_Q, X_i = b, \theta)$; moreover, by definition of max-influence, $\frac{P(X_Q = x_Q | X_i = a, \theta)}{P(X_Q = x_Q | X_i = b, \theta)} \leq e^{e(X_Q | X_i)}$.

Therefore, the right hand side of (7) is at most:

$$e^{e(X_Q|X_i)} \cdot \frac{\int_{x_{R\cup Q}} P(F(X) + \sigma_{\max} Z = w | X_i = a, X_{R\cup Q} = x_{R\cup Q}, \theta) dx_{R\cup Q}}{\int_{x_{R\cup Q}} P(F(X) + \sigma_{\max} Z = w | X_i = b, X_{R\cup Q} = x_{R\cup Q}, \theta) dx_{R\cup Q}}. \tag{8}$$

Since F is 1-Lipschitz, for any fixed value of $X_{R\cup Q}$, $F(X)$ can vary by at most $\texttt{card}(X_N)$ (potentially when all the variables in X_N change values). Since $\sigma_{\max} \geq \frac{\texttt{card}(X_N)}{\epsilon - e(X_Q|X_i)}$ and $Z \sim \text{Lap}(1)$, for any $x_{R\cup Q}$, we have:

$$\frac{P(F(X) + \sigma_{\max} Z = w | X_i = a, X_{R\cup Q} = x_{R\cup Q}, \theta)}{P(F(X) + \sigma_{\max} Z = w | X_i = b, X_{R\cup Q} = x_{R\cup Q}, \theta)} \leq e^{\epsilon - e(X_Q|X_i)}.$$

Combining this with (8), for any w and any secret pair (s_a^i, s_b^i) we have

$$\frac{P(F(X) + \sigma_{\max} Z = w | \theta, s_a^i)}{P(F(X) + \sigma_{\max} Z = w | \theta, s_b^i)} \leq e^{\epsilon},$$

which completes the proof of the theorem. □

B.4 Minimal Markov Quilts Candidates in a Markov Chain

The general Markov Quilt Mechanism in Algorithm 2 searches through all Markov quilts for each X_i. In fact when X is drawn from a Markov chain, and $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$, $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [T]\}$, the mechanism can be made more efficient by only looking at the minimal Markov quilts of X_i. In this section, we formally define minimal Markov Quilt, describe a set of minimal Markov Quilt in the Markov Chain setting in Section 4.1, and prove that it suffices for the Markov Quilt Mechanism to consider only those minimal quilts.

Definition B.1 (Minimal Markov Quilt)**.** *A Markov quilt X_Q (with corresponding X_N, X_R) of X_i is said to be minimal if for any other Markov quilt X'_Q (with corresponding X'_N, X'_R) of X_i such that $X'_N = X_N$, we have $\texttt{card}(X_Q) \leq \texttt{card}(X'_Q)$.*

Lemma B.2 shows a minimal Markov quilt X_Q for each possible X_N. Lemma B.3 shows that it is sufficient to find the optimal noise parameters by only searching through the minimal Markov quilts in Lemma B.2.

Lemma B.2. *Let a, b and i be positive integers such that $i - a \geq 1$ and $i + b \leq T$. Then, $X_Q = \{X_{i-a}, X_{i+b}\}$ is a minimal Markov Quilt for X_i with $X_N = \{X_{i-a+1}, \ldots, X_{i+b-1}\}$, and $X_R = \{X_1, \ldots, X_{i-a-1}\} \cup \{X_{i+b+1}, \ldots, X_T\}$. Moreover, $X_Q = \{X_{i+b}\}$ is a minimal Markov Quilt for X_i with $X_N = \{X_1, \ldots, X_{i+b-1}\}$ and $X_R = \{X_{i+b+1}, \ldots, X_T\}$. Similarly, $X_Q = \{X_{i-a}\}$ is a minimal Markov Quilt for X_i with $X_N = \{X_{i-a+1}, \ldots, X_T\}$ and $X_R = \{X_1, \ldots, X_{i-a-1}\}$. $X_Q = \emptyset$ is a minimal Markov Quilt for X_i with $X_N = \{X_1, \ldots, X_T\}$ and $X_R = \emptyset$.*

Proof. We start with showing the X_Qs in Lemma B.2 are indeed Markov quilts of X_i. First, X_R, X_Q, X_N partitions the node set X, thus meets the first criterion of 4.2. Second, we verify that X_R is conditional independent of X_i given X_Q using the d-separation criteria. If any path from X_i to nodes in X_R is d-separated by X_Q, then the conditional independence claim will hold. In a discrete Markov chain, there exists only one path between X_i and any node $X_j \in X_R$. Since the path connects the nodes in ascending index order, there must be a node in X_Q that blocks the path. Hence X_i is conditionally independent of set X_R given X_Q. The second criterion of 4.2 is met. Therefore, the choices of X_Q in Lemma B.2 are indeed Markov quilts.

It remains to show that X_Q is minimal. First let's consider the case $X_Q = \{X_{i-a}, X_{i+b}\}$ with corresponding $X_N = \{X_{i-a+1}, \ldots, X_{i+b-1}\}$, $i - a \geq 1$ and $i + b \leq T$. Let X'_Q be another quilt such that its corresponding $X'_N = X_N = \{X_{i-a+1}, \ldots, X_{i+b-1}\}$. X_{i-a} must be in X'_Q, otherwise it will be in X'_R and the path between X_{i-a} and X_{i-a+1} cannot be d-separated, which contradicts to the assumption that X'_Q is a Markov quilt. Similarly, X_{i+b} must be in X'_Q, too. Hence $X_Q \subseteq X'_Q$ and X_Q is minimal. In the special case when $X_Q = \{X_{i+b}\}$ and $X_N = \{X_1, \ldots, X_{i+b-1}\}$, X_{i+b} must be in any X'_Q whose corresponding $X'_N = X_N$ because no other node can d-separate X_{i+b} and X_{i+b-1}. The same argument applies to the case when $X_Q = \{X_{i-a}\}$. And obviously for $X_N = \{X_1, \ldots, X_T\}$, there exists only one Markov Quilt which is $X_Q = \emptyset$. Therefore the choices of X_Q in Lemma B.2 are minimal. □

It can be shown that choosing from only those minimal Markov Quilts defined in Lemma B.2 is sufficient.

Lemma B.3. *Let F be a 1-Lipschitz function, and suppose data is represented by a random variable $X = \{X_1, \ldots, X_T\}$ where $X_i \in [k]$ for some positive integer k. Suppose we are given a Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ where $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$, the set of secret pairs $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [T]\}$, and each $\theta \in \Theta$ represents a Markov chain $X_1 \to \cdots \to X_T$.*

For each $i \in [T]$, let $S_{Q,i,mini}$ be the set of minimal Markov Quilts of X_i as defined in Lemma B.2, and $S_{Q,i,all}$ be the set of all Markov Quilts of X_i. Then Markov Quilt Mechanism with $S_{Q,i,mini}$s adds the same amount of noise as Markov Quilt Mechanism with $S_{Q,i,all}$ does.

Proof. We can prove this Lemma by showing that for any Markov Quilt X'_Q (with corresponding X'_N, X'_R) of X_i, its score $\sigma(X_Q)$ as defined in the Markov Quilt Mechanism (Algorithm 2) equals to that of the minimal Markov Quilt $X_Q \in S_{Q,i,\text{mini}}$ (with corresponding X_N, X_R) of X_i if $X_N = X'_N$. Note that the score of a Markov Quilt X_Q (with X_N, X_R) is determined by two parameters: $\texttt{card}(X_N)$ and $e(X_Q|X_i)$. It is obvious that $\texttt{card}(X_N) = \texttt{card}(X'_N)$; so it suffices to prove that $e(X'_Q|X_i) = e(X_Q|X_i)$.

Consider any Markov Quilt X'_Q (with corresponding X'_N, X'_R) of X_i.
If $X_N = \{X_{i-a+1}, \ldots, X_{i+b-1}\}$ for some $i - a \geq 1$ and $i + b \leq T$, i.e., $X_1, X_T \notin X_N$, then X'_Q must be in the form $\{X_{i-a-a'}, \ldots, X_{i-a}, X_{i+b}, \ldots, X_{i+b+b'}\}$ for some $a', b' \geq 0$; and the minimal Markov Quilt $X_Q = \{X_{i-a}, X_{i+b}\}$ is included in $S_{Q,i,\text{mini}}$. Now we compare $e(X'_Q|X_i)$ and $e(X_Q|X_i)$. Let $X'_Q \backslash X_Q = \{X_{i-a-a'}, \ldots, X_{i-a-1}, X_{i+b+1}, \ldots, X_{i+b+b'}\}$. For any $\theta \in \Theta$,

$x, y \in \mathcal{X}$ and $x'_Q \in \mathcal{X}^{\texttt{card}(X'_Q)}$, we have

$$\begin{aligned}
&\frac{P(X'_Q = x'_Q | X_i = x, \theta)}{P(X'_Q = x'_Q | X_i = y, \theta)} \\
=&\frac{P(X_Q = x_Q, X'_Q \backslash X_Q = x'_Q \backslash x_Q | X_i = x, \theta)}{P(X_Q = x_Q, X'_Q \backslash X_Q = x'_Q \backslash x_Q | X_i = y, \theta)} \\
=&\frac{P(X_Q = x_Q | X_i = x, \theta) P(X'_Q \backslash X_Q = x'_Q \backslash x_Q | X_Q = x_Q, \theta)}{P(X_Q = x_Q | X_i = y, \theta) P(X'_Q \backslash X_Q = x'_Q \backslash x_Q | X_Q = x_Q, \theta)} \\
=&\frac{P(X_Q = x_Q | X_i = x, \theta)}{P(X_Q = x_Q | X_i = y, \theta)},
\end{aligned}$$

where the second equality follows from the fact that the set of variables $X'_Q \backslash X_Q$ is independent of X_i given X_Q. Since this equation holds for any $\theta \in \Theta$, $x, y \in \mathcal{X}$ and $x'_Q \in \mathcal{X}^{\texttt{card}(X'_Q)}$, we can conclude that $e(X'_Q | X_i) = e(X_Q | X_i)$.

If $X_N = \{X_1, \ldots, X_{i+b-1}\}$ (or $X_N = \{X_{i-a+1}, \ldots, X_T\}$), we know that any Markov Quilt of X_i must be in the form $X'_Q = \{X_{i+b}, \ldots, X_{i+b+b'}\}$ (or $X'_Q = \{X_{i-a-a'}, \ldots, X_{i-a}\}$) for some $a' \geq 0$ (or $b' \geq 0$); and the minimal Markov Quilt $X_Q = \{X_{i+b}\}$ (or $X_Q = \{X_{i-a}\}$) is included in $S_{Q,i,\text{mini}}$. Again we know that $X'_Q \backslash X_Q$ is independent of X_i given X_Q, and thus can use the argument above to get $e(X'_Q | X_i) = e(X_Q | X_i)$.

If $X_N = \{X_1, \ldots, X_T\}$, the only valid Markov Quilt is $X_Q = \emptyset$, which is included in $S_{Q,i,\text{mini}}$.

Therefore for any X_i and any Markov Quilt X'_Q of it, there exists a Markov Quilt X_Q in $S_{Q,i,\text{mini}}$ such that $\sigma(X'_Q) = \sigma(X_Q)$. So we can conclude that using $S_{Q,i,\text{mini}}$ yields the same result as using $S_{Q,i,\text{all}}$. □

B.5 Max-Influence of X_i on a Markov Quilt X_Q in a Markov Chain

Lemma B.4. *Let X_i be any node in the Markov Chain, and suppose a and b are integers such that $\min(a, b) \geq \frac{\log(1/\pi_\Theta)}{g_\Theta}$. If $X_Q = \{X_{i-a}, X_{i+b}\}$ is a Markov Quilt for X_i, then:*

$$e(X_Q | X_i) \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right) + 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).$$

Moreover, if $X_Q = \{X_{i-a}\}$ is a Markov Quilt for X_i, then:

$$e(X_Q | X_i) \leq 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).$$

and if $X_Q = \{X_{i+b}\}$ is a Markov Quilt for X_i, then:

$$e(X_Q | X_i) \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right).$$

and if $X_Q = \emptyset$ is a trivial Markov Quilt for X_i, then $e(X_Q | X_i) = 0$.

The main ingredient in the proof of Lemma B.4 is the following (fairly standard) result in Markov Chain theory:

Lemma B.5. *Consider a k-state discrete time Markov Chain with a transition matrix P with eigengap g_*. Let π be the stationary distribution of the chain and let $\pi_{\min} = \min_x \pi(x)$ be the minimum probability of any state in the stationary distribution. If P^t is the t-th power of P such that $P^t(x, y) = P(X_{i+t} = y | X_i = x)$, then,*

$$\left|\frac{P^t(x, y)}{\pi(y)} - 1\right| \leq \frac{\exp(-tg_*)}{\pi_{\min}}.$$

Or equivalently,

$$(1 - \Delta_t)\pi(y) \leq P^t(x, y) \leq (1 + \Delta_t)\pi(y), \quad \Delta_t = \frac{\exp(-tg_*)}{\pi_{\min}}.$$

This lemma, along with some algebra, suffices to show the following two facts:

Lemma B.6. *Suppose $t > \frac{\log(1/\pi_\Theta)}{g_\Theta}$, and $\Delta_t = \frac{\exp(-tg_\Theta)}{\pi_\Theta}$. Then, for any j, any $\theta \in \Theta$, and any x, x' and y,*

$$\frac{1-\Delta_t}{1+\Delta_t} \leq \frac{P(X_{t+j} = y | X_j = x, \theta)}{P(X_{t+j} = y | X_j = x', \theta)} \leq \frac{1+\Delta_t}{1-\Delta_t}.$$

Lemma B.7. *Suppose $t > \frac{\log(1/\pi_\Theta)}{g_\Theta}$ and $\Delta_t = \frac{\exp(-tg_\Theta)}{\pi_\Theta}$. Then, for any $j > 0$, any $\theta \in \Theta$, and any x, x' and y,*

$$\left(\frac{1-\Delta_t}{1+\Delta_t}\right)^2 \leq \frac{P(X_j = y | X_{j+t} = x, \theta)}{P(X_j = y | X_{j+t} = x', \theta)} \leq \left(\frac{1+\Delta_t}{1-\Delta_t}\right)^2.$$

B.6 Utility of Markov Quilt Mechanism for Markov Chains

In this section, we prove Theorem 4.4; and based on it, we obtain a consistency lemma for histogram of Markov Chain.

B.6.1 Proof of Theorem 4.4

We restate Theorem 4.4 as Lemma B.8 and proof it.

Lemma B.8 (Utility of Markov Quilt Mechanism for Markov Chains)**.** *Let F be a 1-Lipschitz function, and suppose data is represented by a random variable $X = \{X_1, \ldots, X_T\}$ where $X_i \in [k]$ for some positive integer k. Suppose we are given a Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ where $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$, the set of secret pairs $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [T]\}$, and each $\theta \in \Theta$ represents an irreducible, aperiodic and reversible Markov chain $X_1 \to \cdots \to X_T$ with state space $[k]$.*
Suppose given a realization D of X and the function F, the Markov Quilt Mechanism returns $M(D)$. Let π_Θ and g_Θ be defined as in Section 4.1. If

$$T \geq 2\lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil + 3,$$

then the Laplace noise added by the Markov Quilt Mechanism has scale parameter no larger than C/ϵ for some positive constant C which is determined by Θ and independent of T. Equivalently,

$$\mathbb{E}[|M(D) - F(D)|] \leq \frac{C}{\epsilon}.$$

Proof. Let

$$a = \lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil.$$

Consider the middle node X_i with $i = \lfloor T/2 \rfloor$, and $X_Q = \{X_{i-a}, X_{i+a}\}$, which is guaranteed to be a valid Markov Quilt since $T \geq 2a + 1$. Running Algorithm 5, the process of finding Markov Quilt for each node can stop since neither X_{i-a} nor X_{i+a} is on the boundary of the Markov chain. According to Lemma B.4, we have

$$e(X_Q | X_i) \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right) + 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right) \leq \epsilon/2.$$

Therefore

$$\sigma_{\max} \leq \frac{2a-1}{\epsilon - e(X_Q|X_i)} \leq \frac{4\lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil}{\epsilon},$$

which is independent of T. Since Markov Quilt Mechanism adds Laplace noise with scale parameter $\sigma_{\max}$, we have

$$\mathbb{E}[|M(D) - F(D)|] \leq \frac{4\lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil}{\epsilon}.$$

□

We can compare the utility of Markov Quilt Mechanism with the global sensitivity mechanism for Group DP described in Appendix A.2.1. First, when T is large enough, for any fixed g_Θ and π_Θ, amount of noise does *not* scales with increasing length of the chain T; while the noise added by global sensitivity mechanism always scales linearly with T. This means for large T, the utility of Markov Quilt Mechanism is significantly better than that of global sensitivity mechanism. Also, notice that Markov Quilt Mechanism always considers empty set as one of the Markov Quilt, which, if selected, requires adding the same amount of noise as global sensitivity mechanism does. Therefore we can conclude that Markov Quilt Mechanism is never worse than global sensitivity mechanism for Group DP, and its utility is significantly better especially when T is large.

B.6.2 Consistency for Markov Chains Histogram

Theorem B.9 (Consistency for Markov Quilt Mechanism for Markov Chains Histogram)**.** *Suppose data is represented by a random variable $X = \{X_1, \dots, X_T\}$ where $X_i \in [k]$ for some positive integer k, and suppose the function F returns the histogram density estimator for X. Suppose we are given a Pufferfish framework $(\mathcal{S}, \mathcal{Q}, \Theta)$ where $\mathcal{S} = \{s_a^i : a \in \mathcal{X}, i \in [n]\}$, the set of secret pairs $\mathcal{Q} = \{(s_a^i, s_b^i) : a, b \in \mathcal{X}, a \neq b, i \in [T]\}$, and each $\theta \in \Theta$ represents an irreducible, aperiodic and reversible Markov chain $X_1 \to \cdots \to X_T$ with state space $[k]$.*
Suppose given a realization D of X, the Markov Quilt Mechanism returns $M(D)$ as the private histogram density estimator. Let π_Θ and g_Θ be defined as in Section 4.1. If

$$T \geq 2\lceil \log\left(\frac{\exp(\epsilon/6)+1}{\exp(\epsilon/6)-1}\frac{1}{\pi_\Theta}\right)/g_\Theta \rceil + 3,$$

then we have

$$\mathbb{E}[\|M(D) - F(D)\|_\infty] \leq \frac{C}{\epsilon T}$$

for some positive constant C which is determined by Θ and independent of T.

Proof. Notice that $F(D)$ is a k-dimensional vector and $\|F(D) - F(D')\|_1 \leq 2/T$ for any datasets D, D' that differ by at most one entry. As is mentioned in Appendix B.2, Markov Quilt Mechanism adds Laplace noise to each of the k dimension of $F(D)$. By Lemma B.8, to return a private value of a 1-Lipschitz function F', Markov Quilt Mechanism adds Laplace noise with scale no larger than C/ϵ for some constant C. So it adds Laplace noise with scale parameter no larger than $\frac{2C}{\epsilon T}$ to each dimension of $F(D)$, and we have

$$\mathbb{E}[\|M(D) - F(D)\|_\infty] \leq \frac{2C}{\epsilon T}.$$

□

B.7 Proofs of Lemma B.4, B.5, B.6, B.7

B.7.1 Proof of Lemma B.4

Proof. (of Lemma B.4) First, we can express $e(X_Q|X_i)$ as

$$e(X_Q|X_i) = \max_{x,x' \in \mathcal{X}} \sup_{\theta \in \Theta} \max_{x_Q \in \mathcal{X}^{|Q|}} \log \frac{P(X_Q = x_Q | X_i = x, \theta)}{P(X_Q = x_Q | X_i = x', \theta)}.$$

For any x, x' and any x_Q, by conditional independence,

$$\begin{aligned}\frac{P(X_Q = x_Q|X_i = x, \theta)}{P(X_Q = x_Q|X_i = x', \theta)} =& \frac{P(X_{i+b} = x_{i+b}, X_{i-a} = x_{i-a}|X_i = x, \theta)}{P(X_{i+b} = x_{i+b}, X_{i-a} = x_{i-a}|X_i = x', \theta)} \\ =& \frac{P(X_{i+b} = x_{i+b}|X_i = x, \theta)}{P(X_{i+b} = x_{i+b}|X_i = x', \theta)} \frac{P(X_{i-a} = x_{i-a}|X_i = x, \theta)}{P(X_{i-a} = x_{i-a}|X_i = x', \theta)}.\end{aligned}$$

Taking $\log$ on both sides, we have

$$\begin{aligned}&\log \frac{P(X_Q = x_Q|X_i = x, \theta)}{P(X_Q = x_Q|X_i = x', \theta)} \\ =& \log \frac{P(X_{i+b} = x_{i+b}|X_i = x, \theta)}{P(X_{i+b} = x_{i+b}|X_i = x', \theta)} + \log \frac{P(X_{i-a} = x_{i-a}|X_i = x, \theta)}{P(X_{i-a} = x_{i-a}|X_i = x', \theta)}.\end{aligned}$$

For any $\theta \in \Theta$, let P be the corresponding transition matrix, and π, g_* be the stationary distribution and eigengap of P. Let $\Delta_t^\theta = \frac{\exp(-tg_*)}{\pi_{\min}}$. Then by Lemma B.6 and Lemma B.7, we have

$$\log \frac{P(X_{i+b} = x_{i+b}|X_i = x, \theta)}{P(X_{i+b} = x_{i+b}|X_i = x', \theta)} \leq \log \frac{1+\Delta_t^\theta}{1-\Delta_t^\theta} \leq \log \frac{1+\Delta_t}{1-\Delta_t} = \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right)$$

and

$$\begin{aligned}\log \frac{P(X_{i-a} = x_{i-a}|X_i = x, \theta)}{P(X_{i-a} = x_{i-a}|X_i = x', \theta)} \leq& \log\left(\frac{1+\Delta_t^\theta}{1-\Delta_t^\theta}\right)^2 \leq \log\left(\frac{1+\Delta_t}{1-\Delta_t}\right)^2 \\ =& 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).\end{aligned}$$

Combining the two inequalities together, when $X_Q = \{X_{i-a}, X_{i+b}\}$, we have

$$\log \frac{P(X_Q = x_Q|X_i = x, \theta)}{P(X_Q = x_Q|X_i = x', \theta)} \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right) + 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).$$

Moreover, when $X_Q = \{X_{i-a}\}$, the $\frac{P(X_{i+b}=x_{i+b}|X_i=x,\theta)}{P(X_{i+b}=x_{i+b}|X_i=x',\theta)}$ term will degenerate to 1, thus

$$\log \frac{P(X_Q = x_Q|X_i = x, \theta)}{P(X_Q = x_Q|X_i = x', \theta)} \leq 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right).$$

Similarly, when $X_Q = \{X_{i+b}\}$, the $\frac{P(X_{i-a}=x_{i-a}|X_i=x,\theta)}{P(X_{i-a}=x_{i-a}|X_i=x',\theta)}$ term will degenerate to 1, thus

$$\log \frac{P(X_Q = x_Q|X_i = x, \theta)}{P(X_Q = x_Q|X_i = x', \theta)} \leq \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right).$$

Since the above results hold for any x, x', x_Q and θ, we can conclude the three statements of Lemma B.4. □

B.7.2 Proof to Lemma B.5

The full standard proof can be found at [20].

Proof. Let P be the transition matrix. Construct a matrix A s.t. $A(x,y) = \pi(x)^{1/2}\pi(y)^{-1/2}P(x,y)$. A is symmetric since the chain is reversible, therefore A can be written as

$$A = V\Lambda V^{-1} = V\Lambda V^T,$$

in which V is orthonormal.

Define another inner product

$$\langle \cdot, \cdot \rangle_\pi := \sum_{x\in\Omega} f(x)g(x)\pi(x).$$

We call a matrix π-orthonormal if its basis are orthonormal with respect to $\langle \cdot, \cdot \rangle_\pi$.

Let $D = diag\pi$. Since $A = D^{1/2}PD^{-1/2}$ by the construction of A and $A = V\Lambda V^T$, we have

$$P[D^{-1/2}V] = D^{-1/2}V\Lambda$$

Let $U = D^{-1/2}V$ and let u_i, v_i represent the i-th column vector of U and V respectively, we observe

$$\begin{aligned}\langle u_i, u_j \rangle_\pi =& \sum_{x\in\Omega} u_i(x)u_j(x)\pi(x) \\ =& \sum_{x\in\Omega} u_i(x)\pi(x)^{1/2}u_j(x)\pi(x)^{1/2} \\ =& \langle v_i, v_j \rangle\end{aligned}$$

Thus P is π-orthonormal with $\{u_i\}$ as its basis.

$\langle u_i, u_j \rangle_\pi = 1$ when $i = j$ and $\langle u_i, u_j \rangle_\pi = 0$ otherwise.

Then express P^t in terms of U,

$$\begin{aligned} P^t =& D^{1/2} A^t D^{1/2} \\ =& D^{-1/2} V \Lambda V^T D^{1/2} \\ =& [D^{-1/2} V] \Lambda [V^T D] D^{-1/2} \\ =& U \Lambda U^T D \end{aligned}$$

Therefore,

$$P^t(x, y) = \sum_i u_i(y) \lambda_i^t u_i(x) \pi(y)$$

Since the chain is aperiodic and irreducible, P has a unique largest eigenvalue $\lambda_1 = 1$, we can rewrite the expression above as

$$P^t(x, y) = 1 + \sum_{i=2}^{|\Omega|} u_i(y) \lambda_i^t u_i(x) \pi(y)$$

Apply Cauchy-Schwarz inequality,

$$\begin{aligned} \left| \frac{P^t(x, y)}{\pi(y)} - 1 \right| =& \sum_{i=2}^{|\Omega|} u_i(y) \lambda_i^t u_i(x) \\ \leq& \sum_{i=2}^{|\Omega|} u_i(y) \lambda_*^t u_i(x) \\ \leq& \lambda_*^t \left[\sum_{i=2}^{|\Omega|} u_i(x)^2 \sum_{i=2}^{|\Omega|} u_i(y)^2 \right]^{1/2}, \end{aligned}$$

where λ_* is the eigenvalue of P whose absolute value is the second largest. Furthermore,

$$\sum_{i=2}^{|\Omega|} u_i(x)^2 \leq \sum_{i=1}^{|\Omega|} u_i(x)^2 \leq \pi(x)^{-1}$$

because $\pi(x) = \left\langle \sum_{i=1}^{|\Omega|} u_i(x) \pi(x) u_i, \sum_{i=1}^{|\Omega|} u_i(x) \pi(x) u_i \right\rangle = \pi(x)^2 \sum_{i=1}^{|\Omega|} u_i(x)^2$. We eventually get

$$\left| \frac{P^t(x, y)}{\pi(y)} - 1 \right| \leq \frac{\lambda_*^t}{\sqrt{\pi(x)\pi(y)}} \leq \frac{e^{-(1-\lambda_*)^t}}{\pi_{\min}} = \frac{e^{-g^t}}{\pi_{\min}},$$

where $g = 1 - \lambda_*$ denotes the eigengap of P. □

B.7.3 Proof to Lemma B.6

Proof. Consider any underlying distribution $\theta \in \Theta$ with transition matrix P, and let the stationary distribution be π and the eigengap of P be g_*. We have

$$\frac{P(X_{t+j} = y | X_j = x)}{P(X_{t+j} = y | X_j = x')} = \frac{P^t(x, y)}{P^t(x', y)} = \frac{P^t(x, y)/\pi(y)}{P^t(x', y)/\pi(y)}.$$

Based on Lemma B.5, for $\Delta_t^\theta = \frac{\exp(-t g_*)}{\pi_{\min}}$ with $\pi_{\min} = \min_x \pi(x)$, we have

$$1 - \Delta_t^\theta \leq P^t(x, y)/\pi(y) \leq 1 + \Delta_t^\theta.$$

Recall that $0 \leq \Delta_t^\theta \leq \Delta_t \leq 1$, where Δ_t is as defined in Lemma B.6. Therefore

$$1 - \Delta_t \leq P^t(x, y)/\pi(y) \leq 1 + \Delta_t$$

and the lemma follows. □

B.7.4 Proof to Lemma B.7

Proof. Consider any underlying distribution $\theta \in \Theta$ with transition matrix P, and let the stationary distribution be π and the eigengap of P be g_*. By Bayes' rule we have

$$\begin{aligned}\frac{P(X_j = y|X_{t+j} = x)}{P(X_j = y|X_{t+j} = x')} =& \frac{P(X_{t+j} = x|X_j = y)P(X_j = y)/P(X_{t+j} = x)}{P(X_{t+j} = x'|X_j = y)P(X_j = y)/P(X_{t+j} = x')} \\ =& \frac{P^t(y, x)P(X_{t+j} = x')}{P^t(y, x')P(X_{t+j} = x)}.\end{aligned}$$

Let $\Delta_t^\theta = \frac{\exp(-tg_*)}{\pi_{\min}}$ with $\pi_{\min} = \min_x \pi(x)$. Based on Lemma B.5, we have

$$(1 - \Delta_t^\theta)\pi(x) \leq P^t(y, x) \leq (1 + \Delta_t^\theta)\pi(x).$$

We also have

$$\begin{aligned}P(X_{t+j} = x) &= \sum_y P(X_{t+j} = x|X_j = y)P(X_j = y) \\ &\leq \max_y P(X_{t+j} = x|X_j = y) = \max P^t(y, x) \\ &\leq (1 + \Delta_t^\theta)\pi(x),\end{aligned}$$

and similarly,

$$\begin{aligned}P(X_{t+j} = x) &= \sum_y P(X_{t+j} = x|X_j = y)P(X_j = y) \\ &\geq \min_y P(X_{t+j} = x|X_j = y) = \min_y P^t(y, x) \\ &\geq (1 - \Delta_t^\theta)\pi(x).\end{aligned}$$

Therefore

$$\left(\frac{1 - \Delta_t^\theta}{1 + \Delta_t^\theta}\right)^2 \leq \frac{P^t(y, x)P(X_{t+j} = x')}{P^t(y, x')P(X_{t+j} = x)} \leq \left(\frac{1 + \Delta_t^\theta}{1 - \Delta_t^\theta}\right)^2.$$

The lemma follows from the fact that $0 \leq \Delta_t^\theta \leq \Delta_t \leq 1$, where Δ_t is as defined in Lemma B.6. □

B.8 Implementing the Markov Quilt Mechanism for Markov Chains

In this section, we provide a detailed description of the Markov Quilt Mechanism for Markov chain, along with a more computationally efficient version of the algorithm. Throughout this section we assume that the Pufferfish parameters $(\mathcal{S}, \mathcal{Q}, \Theta)$ and the corresponding g_Θ, π_Θ are known.

B.8.1 The Markov Quilt Mechanism for Markov Chains

Algorithm 3 presents the Markov Quilt mechanism for a Markov Chain. Recall that in this case, the set $S_{Q,i}$ of plausible Markov Quilts for node X_i that we search over is of the form:

$$\begin{aligned}S_{Q,i} =& \Big\{\{X_{i-a}, X_{i+b}\}, \{X_{i-a}\}, \{X_{i+b}\}, \emptyset, \\ & \forall a \in \Big\{\frac{\log(1/\pi_\Theta)}{g_\Theta}, \ldots, i-1\Big\}, b \in \Big\{\frac{\log(1/\pi_\Theta)}{g_\Theta}, \ldots, T-i\Big\}\Big\}. \end{aligned} \tag{9}$$

Algorithm 3 enumerates over all Markov Quilts in $S_{Q,i}$ for all nodes X_i in the chain, computes the requisite scores, and calculates $\sigma_{\max}$, which is (approximately) the minimum standard deviation of the amount of noise needed to keep the values of all nodes private. Finally, it adds Laplace noise with standard deviation proportional to $\sigma_{\max}$ to the exact value of $F(D)$ to ensure privacy.

Algorithm 4 shows the FindBestMarkovQuilt function that returns the minimal score of any Markov Quilt in $S_{Q,i}$. This is computed by searching over the set $S_{Q,i}$ of candidate Markov quilts for X_i.

Algorithm 3 BasicMarkovQuiltMechansim(ϵ, F, D)

Initialize $\sigma_{\max} \leftarrow 0$
for $X_i \in X$ **do**
  $\sigma_i \leftarrow$ FindBestMarkovQuilt(i, ϵ)
  $\sigma_{\max} = \max(\sigma_{\max}, \sigma_i)$
end for
return $F(D) + \sigma_{\max} \cdot \text{Lap}(1)$

Algorithm 4 FindBestMarkovQuilt(i, ϵ)

$\sigma_i \leftarrow \infty$
$(a^*, b^*) \leftarrow (i, T - i + 1)$
for $(X_{i-a}, X_{i+b}), a, b \geq \log(1/\pi_\Theta)/g_\Theta$ that is a Markov Quilt for X_i **do**
  if $i - a = 0$ and $i + b = T + 1$ **then**
    $e(a, b) \leftarrow 0$
  else if $i - a = 0$ **then**
    $e(a, b) \leftarrow \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right)$
  else if $i + a = T + 1$ **then**
    $e(a, b) \leftarrow 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right)$
  else
    $e(a, b) \leftarrow \log\left(\frac{\pi_\Theta + \exp(-g_\Theta b)}{\pi_\Theta - \exp(-g_\Theta b)}\right) + 2\log\left(\frac{\pi_\Theta + \exp(-g_\Theta a)}{\pi_\Theta - \exp(-g_\Theta a)}\right)$
  end if
  if $e(a, b) > \epsilon$ **then**
    Pass
  else
    $\sigma_i \leftarrow \min(\frac{a+b-1}{\epsilon - e(a,b)}, \sigma_i)$
    If $\sigma_i = \frac{a+b-1}{\epsilon - e(a,b)}$ then $(a^*, b^*) \leftarrow (a, b)$.
  end if
end for
return σ_i, a^*, b^*

B.8.2 A Faster Algorithm

BasicMarkovQuiltMechansim needs to enumerate over all nodes in the chain and compute the best Markov quilt for each one. This can be improved in most cases. The intuition is that the best Markov quilt for a node near the center of the chain should also be a privacy-preserving choice of quilt for a node near the end of the chain. For ease of presentation, we introduce two imaginary nodes X_0 and X_{T+1} at either end of the chain. We prepend X_0 to to X_1 and append X_{T+1} to X_n. Quilt $\{X_0, X_{i+b}\}$ in the algorihm description corresponds to Quilt $\{X_{i+b}\}$ in $S_{Q,i}$, $\{X_{i-a}, X_{T+1}\}$ corresponds to Quilt $\{X_{i-a}\}$, and Quilt $\{X_0, X_{T+1}\}$ corresponds to the trivial quilt $\emptyset$.

We shall show that the noise parameter σ calculated from this quilt is enough to work for all nodes, that is, preserves ϵ-Pufferfish privacy for all nodes. Starting from a node in the middle of the chain offers better chance of finding such a quilt than starting from the end. Algorithm 5 shows the algorithm.

Claim B.10. *The standard deviation of the noise added to $F(D)$ by BasicMarkovQuiltMechansim is larger than or equal to that added by FastMarkovQuiltMechanism.*

Proof. This is obvious. The $\sigma_{\max}$ in BasicMarkovQuiltMechansim is the maximum of FindBestMarkovQuilt(i, ϵ) for all X_is in $\mathcal{X}$, while FastMarkovQuiltMechanism's $\sigma_{\max}$ is the maximum of of FindBestMarkovQuilt(i, ϵ) for only a subset of $\mathcal{X}$. Therefore, the amount of noise added to $F(D)$ by BasicMarkovQuiltMechansim cannot be smaller. □

Claim B.11. *Suppose the best Markov Quilt found by FastMarkovQuiltMechanism is $\{X_{i-a}, X_{i+b}\}$ where $i - a \geq 1$ and $i + b \leq T$. Then, FastMarkovQuiltMechanism ensures ϵ-Pufferfish privacy.*

Algorithm 5 FastMarkovQuiltMechanism(ϵ, F, D)

Initialize $\sigma_{\max} \leftarrow 0$
$T = |\mathcal{X}|$
$i = \lfloor \frac{T}{2} \rfloor$
while i has not being explored **do**
  Mark i as explored
  $(a_i, b_i, \sigma_i) \leftarrow$ FindBestMarkovQuilt (i, ϵ)
  $\sigma_{\max} = \max(\sigma_{\max}, \sigma_i)$
  if $i - a_i = 0$ and $i + b_i = T + 1$ **then**
    break
  else if $i - a_i = 0$ **then**
    $i = i + 1$
  else if $i + b_i = T + 1$ **then**
    $i = i - 1$
  else
    break
  end if
end while
return $F(D) + \sigma_{\max} \cdot \text{Lap}(1)$

Proof. First let $j \neq i$ be such that $j - a_i \geq 1$ and $j + b_i \leq T$. Then, $\{X_{j-a}, X_{j+b}\}$ is a Markov Quilt for X_j with score equal to σ_i, and thus $\sigma_{\max} = \sigma_i \geq \sigma_j$. Thus the amount of noise added by FastMarkovQuiltMechanism is sufficient to protect the privacy of X_j.

Next consider a $j \neq i$ such that $j - a_i \geq 1$ but $j + b_i > T$. Then, $\{X_{j-a}\}$ is a Markov Quilt for X_j with score less than σ_i. Thus $\sigma_{\max} = \sigma_i \geq \sigma_j$ in this case as well. The final case, where $j - a_i < 1$ but $j + b_i \leq T$ is analogous. Thus, the amount of noise added by FastMarkovQuiltMechanism is sufficient to protect the privacy of such X_j's, and as such it preserves ϵ-Pufferfish privacy. □

The rest of the proof follows from the observation that FastMarkovQuiltMechanism either finds a node X_i and a corresponding Markov Quilt $\{X_{i-a}, X_{i+b}\}$ such that $i - a \geq 1$ and $i + b \leq T$; or, finds the empty Markov Quilt (in which case it adds noise with standard deviation T/ϵ.

B.9 More about Simulation

B.9.1 Binary Markov Chain

Lemma B.12. *Any binary Markov Chain with transition matrix P such that $P_{ij} > 0$ for all $i, j \in \{1, 2\}$ is* aperiodic*,* irreducible *and* reversible*.*

Proof. Let $p_0 = P(X_{i+1} = 0 | X_i = 0)$ and $p_1 = P(X_{i+1} = 1 | X_i = 1)$, both lies in $(0, 1)$. So we have

$$P = \begin{bmatrix} p_0 & 1 - p_0 \\ 1 - p_1 & p_1 \end{bmatrix}.$$

The stationary distribution of the chain is

$$[\pi_0, \pi_1] = [\frac{1 - p_1}{2 - p_1 - p_0}, \frac{1 - p_0}{2 - p_1 - p_0}].$$

So we have

$$\pi_0 P(X_{i+1} = 1 | X_i = 0) = \frac{(1 - p_1)(1 - p_0)}{2 - p_1 - p_0}$$
$$\pi_1 P(X_{i+1} = 0 | X_i = 1) = \frac{(1 - p_0)(1 - p_1)}{2 - p_1 - p_0},$$

which means the chain is *reversible*.
Since every element in P is positive, we know the chain is *aperiodic* and *irreducible*. □

B.10 Parameters in Inferential Privacy [11]

[11] defines multiplicative influence between any two variables X_i, X_j as

$$\gamma_{ij} = \max\left\{\frac{1}{2}\log\frac{P(X_i \in S|X_{-i} = x_{-i})}{P(X_i \in S|X_{-i} = x'_{-i})} \;\middle|\; S \subseteq \text{supp}(X_i), x_{-i} \sim_j x'_{-i}\right\},$$

where $X_{-i} = (X_1, \ldots, X_{i-1}, X_{i+1}, \ldots, X_n)$, and $x_{-i} \sim_j x'_{-i}$ means that x_{-i}, x_{-i} differ only on value of X_j. Using a convention that $\gamma_{ii} = 0$, all γ_{ij}s together form an $n \times n$ multiplicative influence matrix $\Gamma = (\gamma_{ij})$. Let $\Phi = (\phi_{ij}) = (I - \Gamma)^{-1}$. [11] shows that if the spectral norm of Γ is less than 1, then an algorithm that guarantees ϵ-differential privacy provides $2\epsilon\sum_{j=1}^{n}\phi_{ij}$ inferential privacy for all X_i.

In the binary Markov Chain setting, since a node is independent of every other node given its two neighbors, Γ is all zero except for sub-diagonal entries. Let $p_0 = P(X_{i+1} = 0|X_i = 0)$ and $p = P(X_{i+1} = 1|X_i = 1)$. It can be shown that Γ is a matrix such that $\gamma_{ij} = 0$ for all $i = j, |i - j| > 1$; and for all $|i - j| = 1$,

$$\begin{aligned}
&\gamma_{ij} = \gamma = \frac{1}{2}\max\Bigg\{\\
&\pm\log\left(\frac{1-p_1}{1-p_0}\frac{(1-p_0)^2 + (1-p_1)p_0}{(1-p_1)(1-p_0) + p_1(1-p_1)}\right), \pm\log\left(\frac{p_1}{p_0}\frac{(1-p_0)^2 + (1-p_1)p_0}{(1-p_1)(1-p_0) + p_1(1-p_1)}\right),\\
&\pm\log\left(\frac{1-p_1}{1-p_0}\frac{p_0(1-p_0) + p_0p_1}{(1-p_1)p_0 + p_1^2}\right), \pm\log\left(\frac{p_1}{p_0}\frac{p_0(1-p_0) + p_0q_0}{(1-p_1)p_0 + p_1^2}\right)\Bigg\}.
\end{aligned}$$

We have

$$\Gamma = \begin{bmatrix}
0 & \gamma & 0 & 0 & \cdots & 0\\
\gamma & 0 & \gamma & 0 & \cdots & 0\\
0 & \gamma & 0 & \gamma & \cdots & 0\\
 & & \cdots & & & \\
0 & \cdots & 0 & \gamma & 0 & \gamma\\
0 & \cdots & 0 & 0 & \gamma & 0
\end{bmatrix}.$$

Notice that the algorithm requires that the spectral norm of Γ is less than 1. This condition fails for some p, q; and the region where it is satisfied also depends on the size of the matrix, i.e., length of the chain. It can be shown that as the length of the chain $n \to \infty$, we need $\gamma < 1/2$, because

$$\Gamma\begin{bmatrix}1\\1\\\cdots\\1\\1\end{bmatrix} = \begin{bmatrix}\gamma\\2\gamma\\\cdots\\2\gamma\\\gamma\end{bmatrix},$$

and

$$\lim_{n\to\infty}\frac{\|[\gamma, 2\gamma, \ldots, 2\gamma, \gamma]^\top\|_2}{\|[1, \ldots, 1]^\top\|_2} = 2\gamma.$$

If the case where the spectral norm of Γ is less than 1, we compute the amount of noise in the following way. Since we require the algorithm to guarantee ϵ-Pufferfish Privacy for all X_i, we set $2\epsilon'\max_i\sum_{j=1}^{n}\phi_{ij} = \epsilon$, and add noise $Z \sim \text{Lap}(1/\epsilon')$ to the true value $F(D)$.

B.11 More Comparison with Inferential Privacy [11]

In this section, we present a more detailed comparison between Markov Quilt Mechanism and Inferential Privacy [11], with data from binary Markov chain.

The setting is the same as that in Section 5. We consider a 1-Lipschitz function F that outputs the summation S of states over time $t = 1, \ldots, T$. We represent each distribution θ with two parameters: $p_0 = P(X_{i+1} = 0|X_i = 0)$ and $p_1 = P(X_{i+1} = 1|X_i = 1)$. Given a pair of (p_0, p_1), we examine the scale of the Laplace noise the two algorithms add.

Specifically, we fix the length of the chain to be $T = 100$, vary $p_0, p_1 \in \{0.01, 0.02, \ldots, 0.99\}$, and

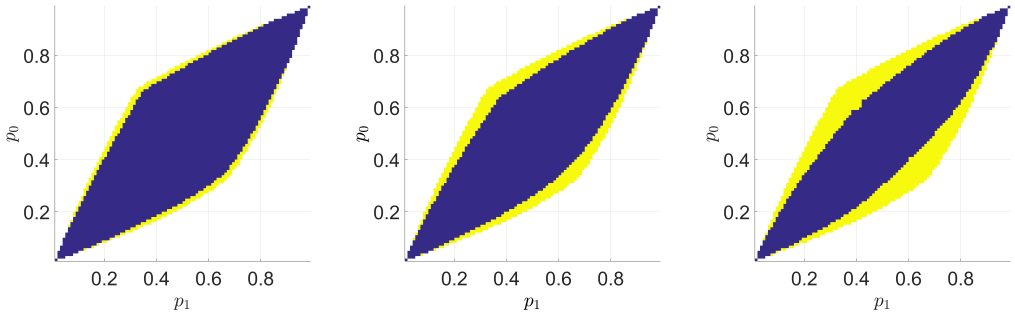

(a) $\epsilon = 1$. (b) $\epsilon = 5$. (c) $\epsilon = 10$.

Figure 3: White region is where Inferential Privacy fails to provide guarantee. Yellow region is where Markov Quilt Mechanism adds less noise. Blue region is where Inferential Privacy adds less noise.

$\epsilon \in \{1, 5, 10\}$. For each ϵ, we compare the noise added by both algorithms. Figure 3 shows the region where one algorithm outperform the other. We notice that there is a pretty large region where Inferential Privacy fails to provide any guarantee. And we can see that as ϵ grows, Markov Quilt Mechanism performs increasingly better than Inferential Privacy, which is because Inferential Privacy adds noise which is linear in $1/\epsilon$, yet the noise scale of Markov Quilt Mechanism decreases faster.

B.12 More Simulations for Markov Quilt Mechanism

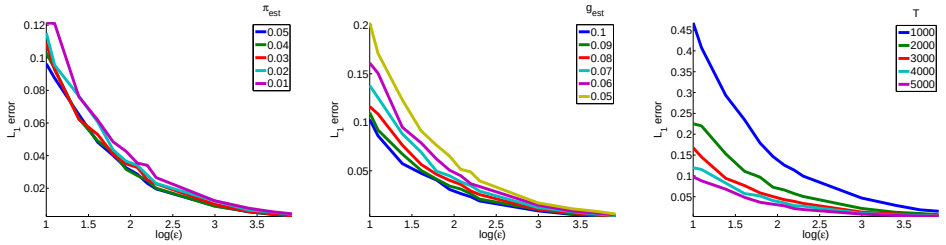

(a) Vary π_{est}. ($g_{\text{est}} = 0.1$, $T = 5000$). (b) Vary g_{est}. ($\pi_{\text{est}} = 0.05$, $T = 5000$). (c) Vary T. ($g_{\text{est}} = 0.1$, $\pi_{\text{est}} = 0.05$).

Figure 4: Averaged L_1 distance between private and non-private histograms vs. $\log(\epsilon)$. Different lines represent different values of (a).π_{est} (b).g_{est} (c). T.

In this section, we conduct more simulations to illustrate the performance of the Markov Quilt Mechanism. We consider the setting in Section 4.1 where the underlying Bayesian network is an *aperiodic*, *irreducible* and *reversible* discrete time Markov chain with k states. The function F we use is the histogram of states over time $t = 1, \ldots, T$. The output of F is a k-dimensional vector $(f_1, \ldots, f_k)$ where f_i is the frequency of state i in times 1 to T, and the desired output is noisy frequencies $(f'_1, \ldots, f'_k)$. This histogram function is $2/T$-Lipschitz with respect to the L_1 norm, and thus can be scaled to fit our mechanism, as is described in Appendix B.2.

Synthetic data is generated as follows. Firstly, we choose two generating parameters $\pi_{\text{gen}} \in (0, 1/k)$, $g_{\text{gen}} \in (0, 1)$ and generate an *aperiodic*, *irreducible* and *reversible* transition matrix P such that $\pi_{\min}(P) \geq \pi_{\text{gen}}$ and $g(P) \geq g_{\text{gen}}$. Then we generate a random vector $q \in \mathbb{R}^k$ uniformly from the probability simplex as the initial distribution. A sequence $X = \{X_i\}_{i=1}^T, X_i \in [k]$ is then generated from the Markov Chain induced by P and q as our dataset.

Notice that the transition matrix P or the generating parameters $\pi_{\text{gen}}, g_{\text{gen}}$ are not known to the private mechanism; the mechanism only knows conservative estimates of π_{gen} and g_{gen}, denoted by π_{est} and g_{est}. We enforce that $\pi_{\text{est}} \leq \pi_{\text{gen}}$ and $g_{\text{est}} \leq g_{\text{gen}}$. Therefore, $\pi_{\text{est}}, g_{\text{est}}$ together with ϵ can be viewed as a set of "privacy parameters". As ϵ gets smaller, the mechanism guarantees more privacy; as π_{est} or g_{est} gets smaller, the mechanism is able to guarantee protection against a larger set of distributions.

Utility of the mechanism is measured by the L_1 distance between the output $(f'_1, \ldots, f'_k)$ with $(f_1, \ldots, f_k)$, the exact frequencies. A smaller distance implies higher utility. In our simulations, we investigate the effect on utility of the three privacy parameters, as well as T, the total length of the sequence. In each simulation, we keep two of the parameters $\{\pi_{\text{est}}, g_{\text{est}}, T\}$ fixed, and vary the remaining one and ϵ. We plot L_1 distance on the y-axis as a function of $\log \epsilon$. P, q are generated with $k = 5$ and generating parameters $\pi_{\text{gen}} = 0.05, g_{\text{gen}} = 0.1$, and are fixed throughout the simulations. We repeat each simulation 100 times and plot the averaged error in the figures. A new dataset $X = \{X_i\}_{i=1}^{T}$ is generated at each repeated run.

Figure 4 shows the results. We observe that as expected, when the remaining parameters are fixed, utility improves with larger ϵ as well as larger π_{est} and larger g_{est}. An interesting observation is that π_{est} is less influential to utility than g_{est}; this agrees with the dependence of max-influence on π_{est} in Lemma B.6. Finally, Figure 4c shows that as the total sequence length T gets larger, our Markov Quilt Mechanism outputs more accurate results.